

1'MM 82 518

(12) UK Patent Application (19) GB (11) 2 390 775 (13) A

(43) Date of A Publication

14.01.2004

(21) Application No:	0216245.1	(51) INT CL ⁷ : H04L 12/56 // H04Q 7/22
(22) Date of Filing:	12.07.2002	(52) UK CL (Edition W): H4K KTKX H4L LDGP L201 L203 L205 L213
(71) Applicant(s): Fujitsu Limited (Incorporated in Japan) 1-1 Kamikodanaka 4-chome, Nakahara-ku, Kawasaki-shi, Kanagawa 211-8588, Japan		(56) Documents Cited: GB 2361147 A GB 2355623 A GB 2313264 A GB 2299729 A
(72) Inventor(s): Abedi Saled		(58) Field of Search: UK CL (Edition V) H4K KTKA KTKX, H4L LDGP INT CL ⁷ H04L 12/56 29/06, H04Q 7/22 Other: Online: WPI, EPODOC, JAPIO and the Internet
(74) Agent and/or Address for Service: Haseltine Lake & Co Imperial House, 15-19 Kingsway, LONDON, WC2B 6UD, United Kingdom		

(54) Abstract Title: **Packet scheduler using a genetic algorithm**

(57) Packet scheduling apparatus 10 schedules packets of data for transmission from a transmitter via at least one channel to a plurality of receivers.

A candidate solution generating unit 12 generates a plurality of candidate scheduling solutions $\eta_1 \dots \eta_n$. At least one candidate solution is generated using a genetic algorithm. Each candidate solution specifies at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration.

A best solution selecting unit 14 compares the generated candidate solutions and, based on the comparison results, selects a best one of the candidate solutions η_{best} to use to transmit packets in said scheduling instant.

Such apparatus can provide a superior scheduling solution to conventional schedulers such as round-robin and maximum carrier-to-interference ratio schedulers whilst keeping the computation resource requirements within acceptable levels.

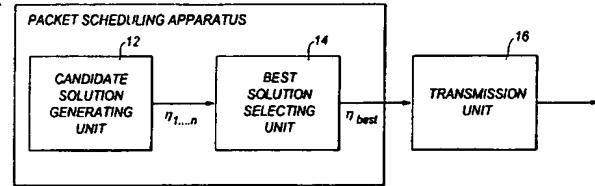


Fig.3

GB 2 390 775 A

MISSOURI BANK (1850)

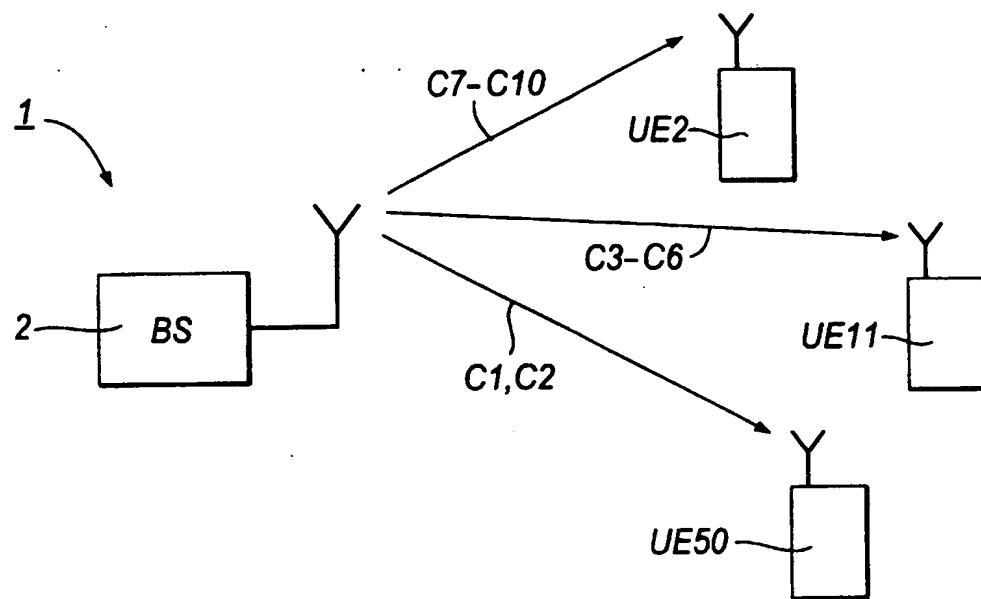


Fig. 1

		TTI 1	TTI 2	TTI 3	TTI 4	TTI 5	TTI 6	TTI 7	TTI 8
Node B 1 Channel Code 1	UE50 Packet 1	UE1 Packet 1	UE3 Packet 1	UE23 Packet 1	UE4 Packet 2	UE6 Packet 1	UE1 Packet 1	UE23 Packet 1	UE7 Packet 1
	UE50 Packet 2	UE50 Packet 3	UE4 Packet 1	UE1 Packet 2	UE7 Packet 1	UE50 Packet 2	UE4 Packet 1	UE1 Packet 2	UE17 Packet 1
Node B 1 Channel Code 2	UE11 Packet 1	UE50 Packet 4	UE1 Packet 1	UE4 Packet 2	UE23 Packet 2	UE2 Packet 8	UE11 Packet 4	UE15 Packet 1	UE23 Packet 2
	UE11 Packet 2	UE50 Packet 5	UE50 Packet 10	UE11 Packet 8	UE16 Packet 2	UE9 Packet 2	UE40 Packet 2	UE34 Packet 1	UE9 Packet 3
Node B 1 Channel Code 3	UE11 Packet 3	UE2 Packet 5	UE16 Packet 1	UE9 Packet 1	UE11 Packet 3	UE24 Packet 2	UE43 Packet 2	UE16 Packet 1	UE11 Packet 10
	UE11 Packet 4	UE11 Packet 5	UE11 Packet 6	UE2 Packet 3	UE1 Packet 3	UE4 Packet 1	UE11 Packet 9	UE40 Packet 3	UE11 Packet 5
Node B 1 Channel Code 4	UE11 Packet 5	UE2 Packet 3	UE16 Packet 1	UE9 Packet 1	UE11 Packet 3	UE24 Packet 2	UE43 Packet 2	UE16 Packet 1	UE11 Packet 10
	UE11 Packet 6	UE11 Packet 7	UE11 Packet 8	UE2 Packet 6	UE1 Packet 6	UE4 Packet 1	UE11 Packet 9	UE40 Packet 3	UE11 Packet 5
Node B 1 Channel Code 5	UE11 Packet 7	UE2 Packet 4	UE50 Packet 6	UE3 Packet 11	UE24 Packet 1	UE2 Packet 8	UE38 Packet 1	UE38 Packet 5	UE50 Packet 14
	UE11 Packet 8	UE2 Packet 5	UE50 Packet 7	UE11 Packet 9	UE43 Packet 1	UE4 Packet 3	UE38 Packet 2	UE38 Packet 2	UE14 Packet 1
Node B 1 Channel Code 6	UE11 Packet 9	UE2 Packet 6	UE50 Packet 8	UE3 Packet 12	UE24 Packet 9	UE2 Packet 1	UE38 Packet 1	UE38 Packet 5	UE50 Packet 14
	UE11 Packet 10	UE2 Packet 7	UE50 Packet 9	UE11 Packet 12	UE43 Packet 9	UE4 Packet 3	UE38 Packet 2	UE38 Packet 2	UE14 Packet 1
Node B 1 Channel Code 7	UE11 Packet 11	UE2 Packet 8	UE50 Packet 10	UE3 Packet 13	UE24 Packet 1	UE2 Packet 8	UE38 Packet 1	UE38 Packet 5	UE50 Packet 14
	UE11 Packet 12	UE2 Packet 9	UE50 Packet 11	UE11 Packet 13	UE43 Packet 1	UE4 Packet 3	UE38 Packet 2	UE38 Packet 2	UE14 Packet 1
Node B 1 Channel Code 8	UE11 Packet 13	UE2 Packet 10	UE50 Packet 12	UE3 Packet 14	UE24 Packet 1	UE2 Packet 9	UE38 Packet 1	UE38 Packet 5	UE50 Packet 14
	UE11 Packet 14	UE2 Packet 11	UE50 Packet 13	UE11 Packet 15	UE43 Packet 1	UE3 Packet 3	UE38 Packet 1	UE38 Packet 3	UE14 Packet 1
Node B 1 Channel Code 9	UE11 Packet 15	UE2 Packet 12	UE50 Packet 14	UE3 Packet 16	UE24 Packet 1	UE2 Packet 10	UE38 Packet 1	UE38 Packet 3	UE43 Packet 2
	UE11 Packet 16	UE2 Packet 13	UE50 Packet 15	UE11 Packet 17	UE43 Packet 1	UE2 Packet 11	UE40 Packet 3	UE40 Packet 3	UE43 Packet 2
Node B 1 Channel Code 10	UE11 Packet 17	UE2 Packet 14	UE50 Packet 16	UE3 Packet 18	UE24 Packet 1	UE2 Packet 12	UE38 Packet 4	UE38 Packet 4	UE43 Packet 2
	UE11 Packet 18	UE2 Packet 15	UE50 Packet 17	UE11 Packet 19	UE43 Packet 1	UE2 Packet 13	UE50 Packet 4	UE50 Packet 4	UE43 Packet 2

Fig. 2

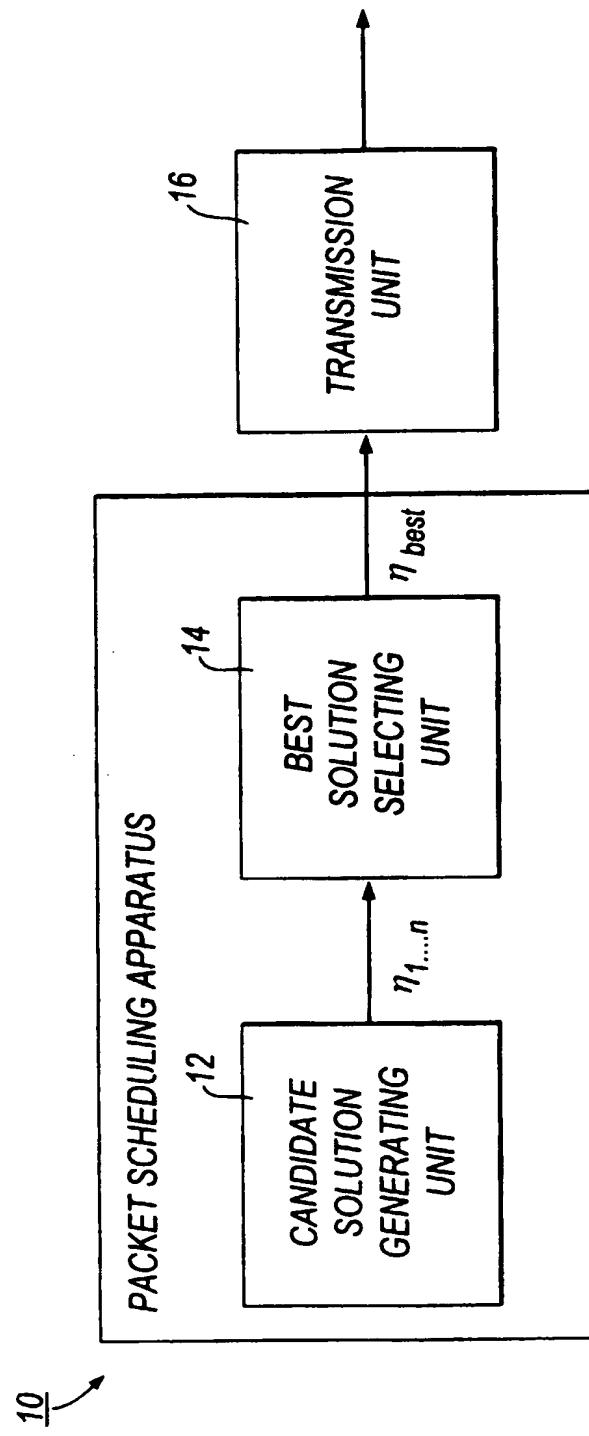


Fig. 3

4/22

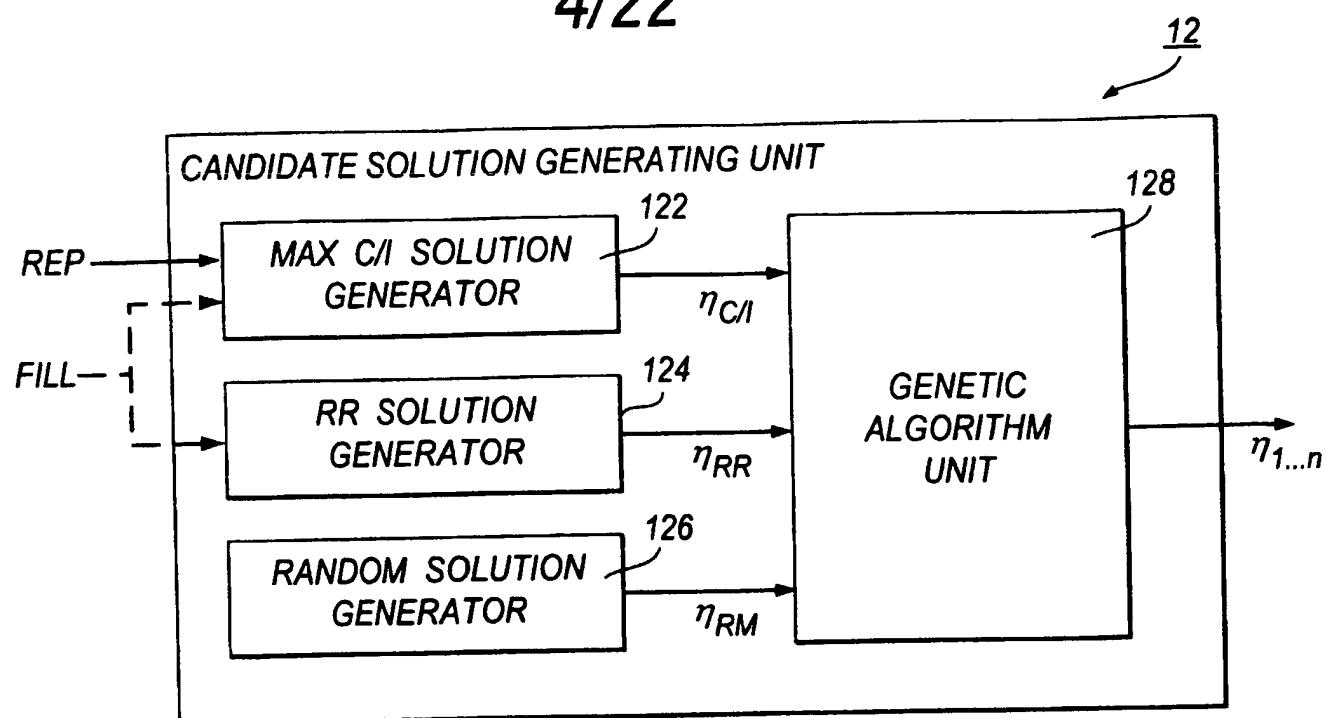


Fig.4

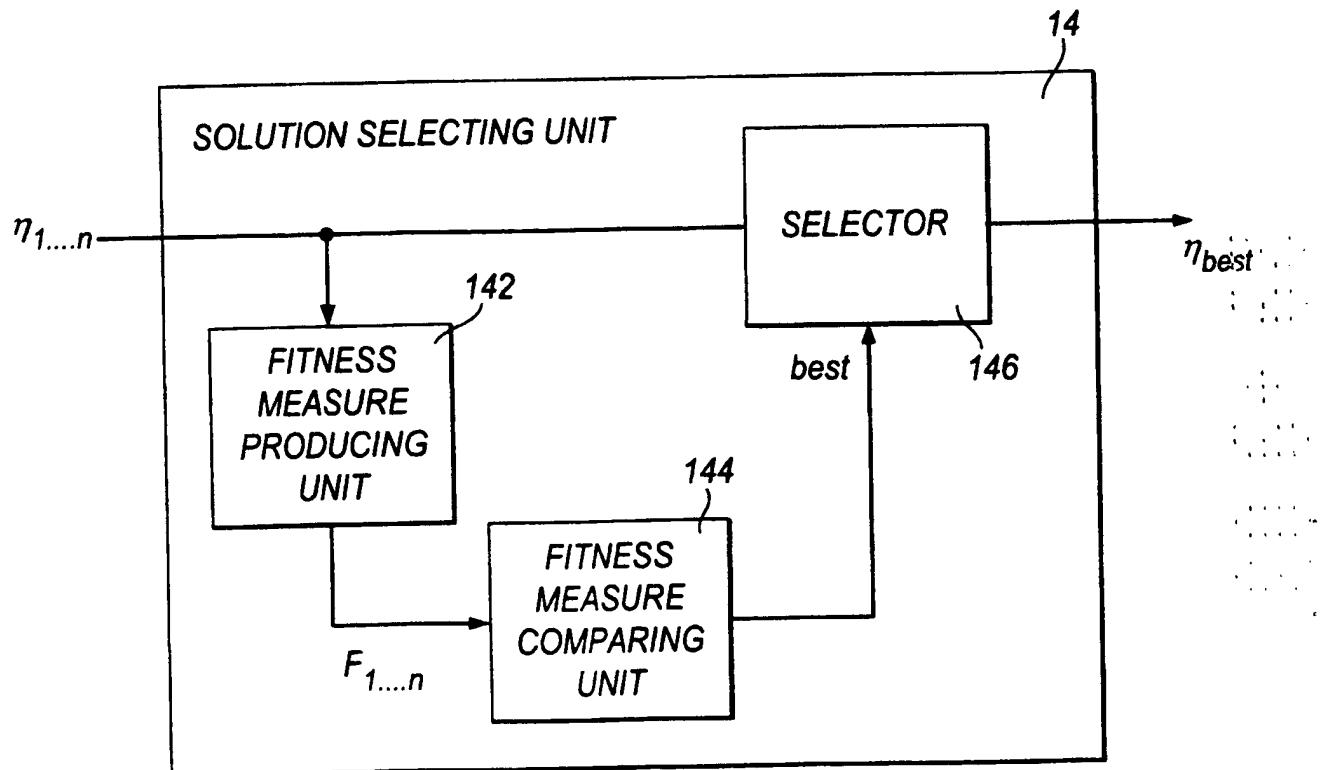


Fig.5

5/22

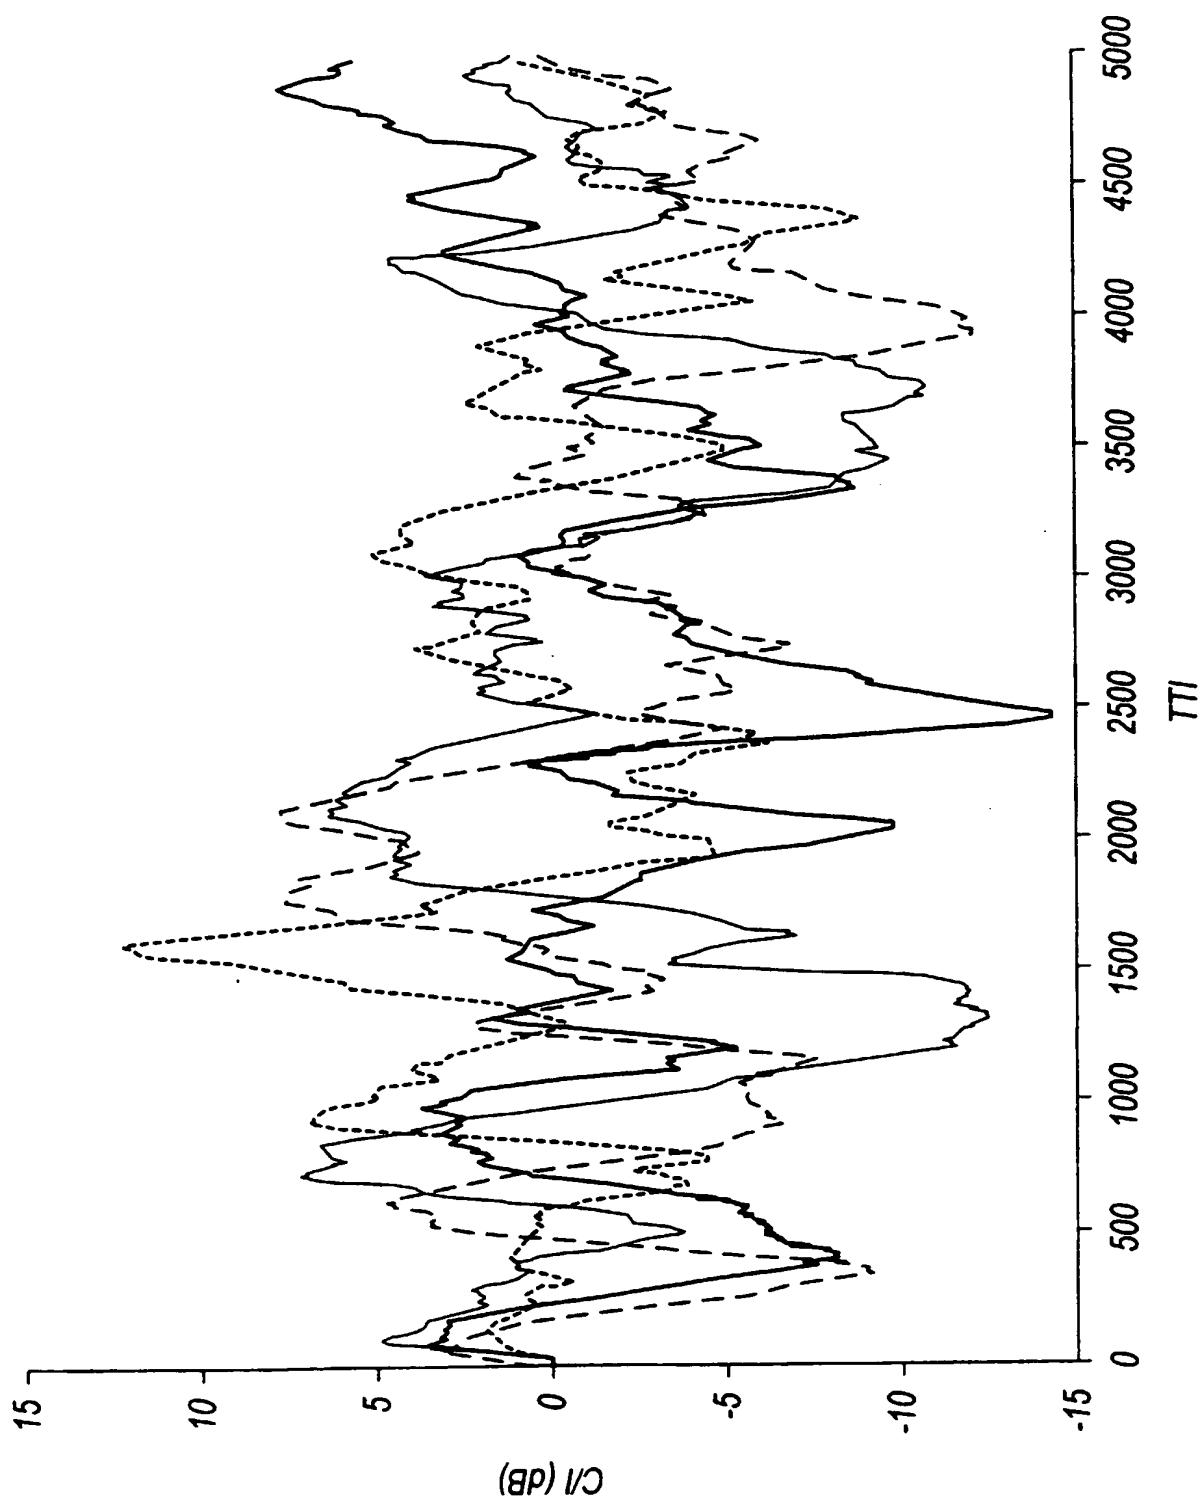


Fig. 6

6/22

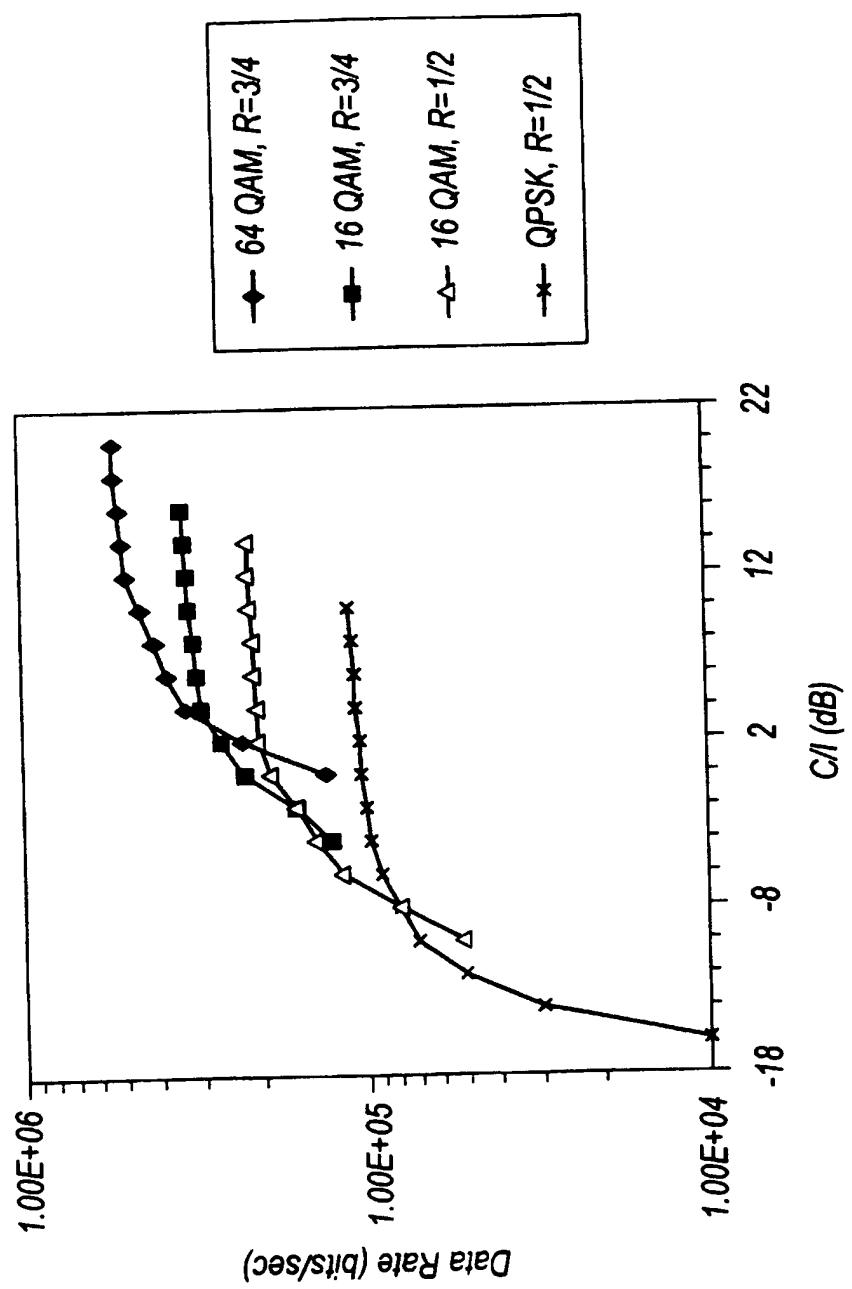


Fig. 7

7/22

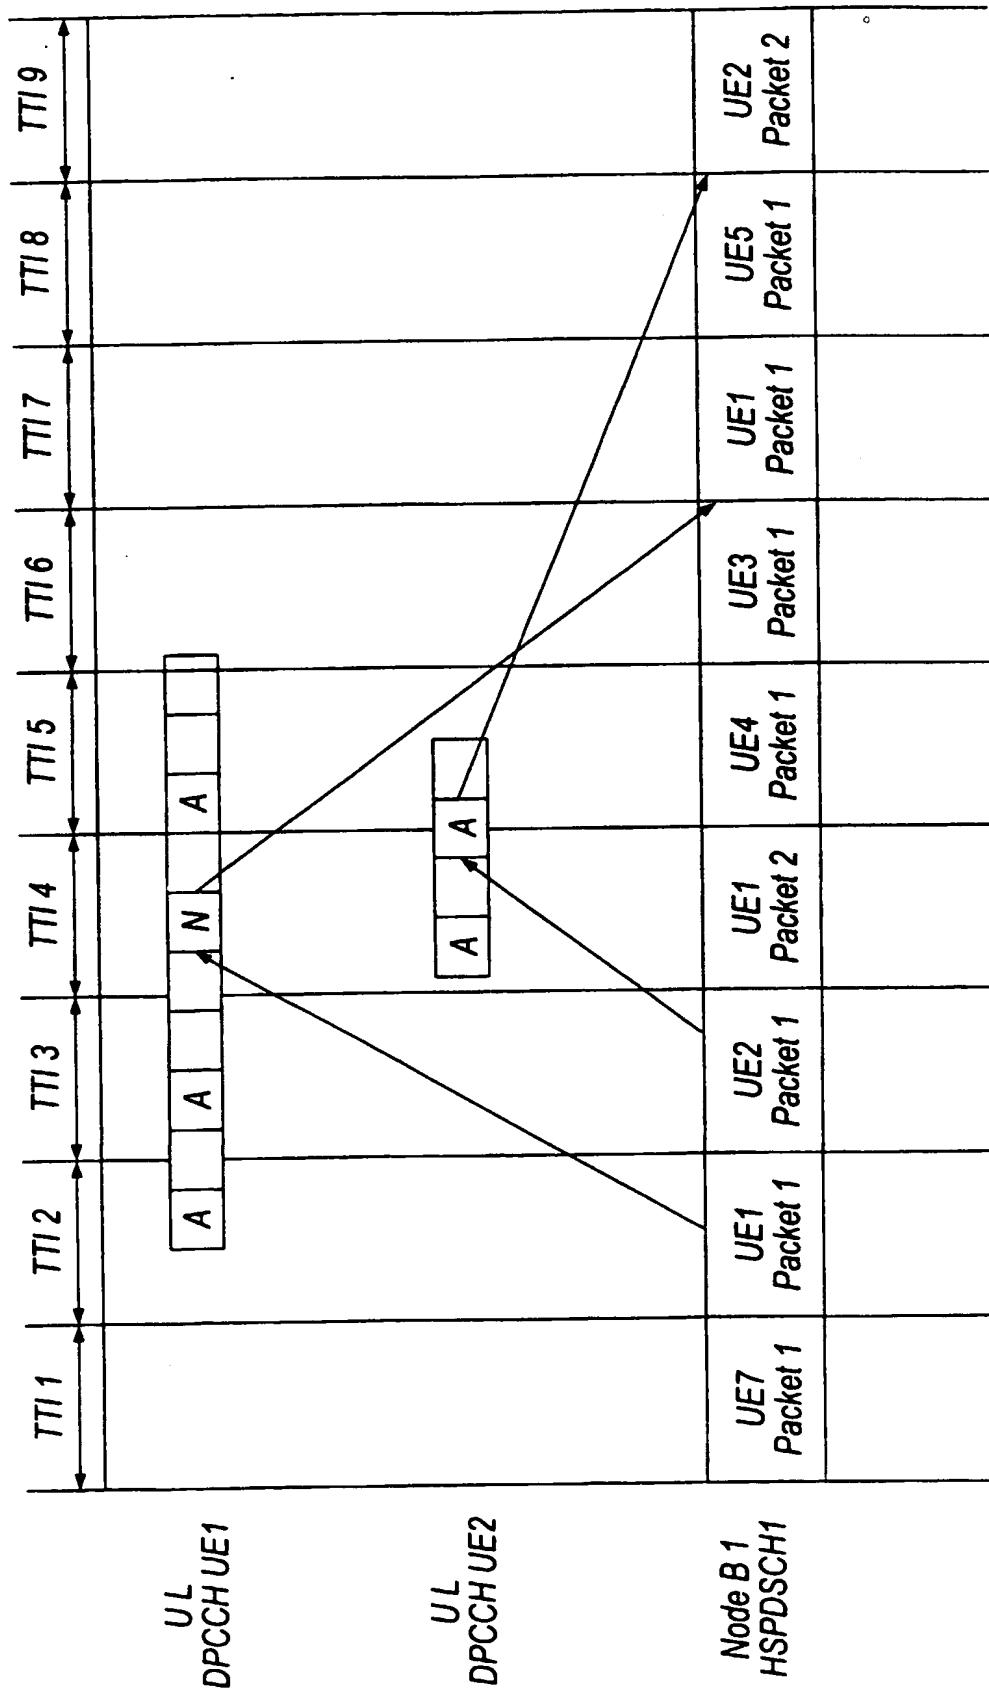
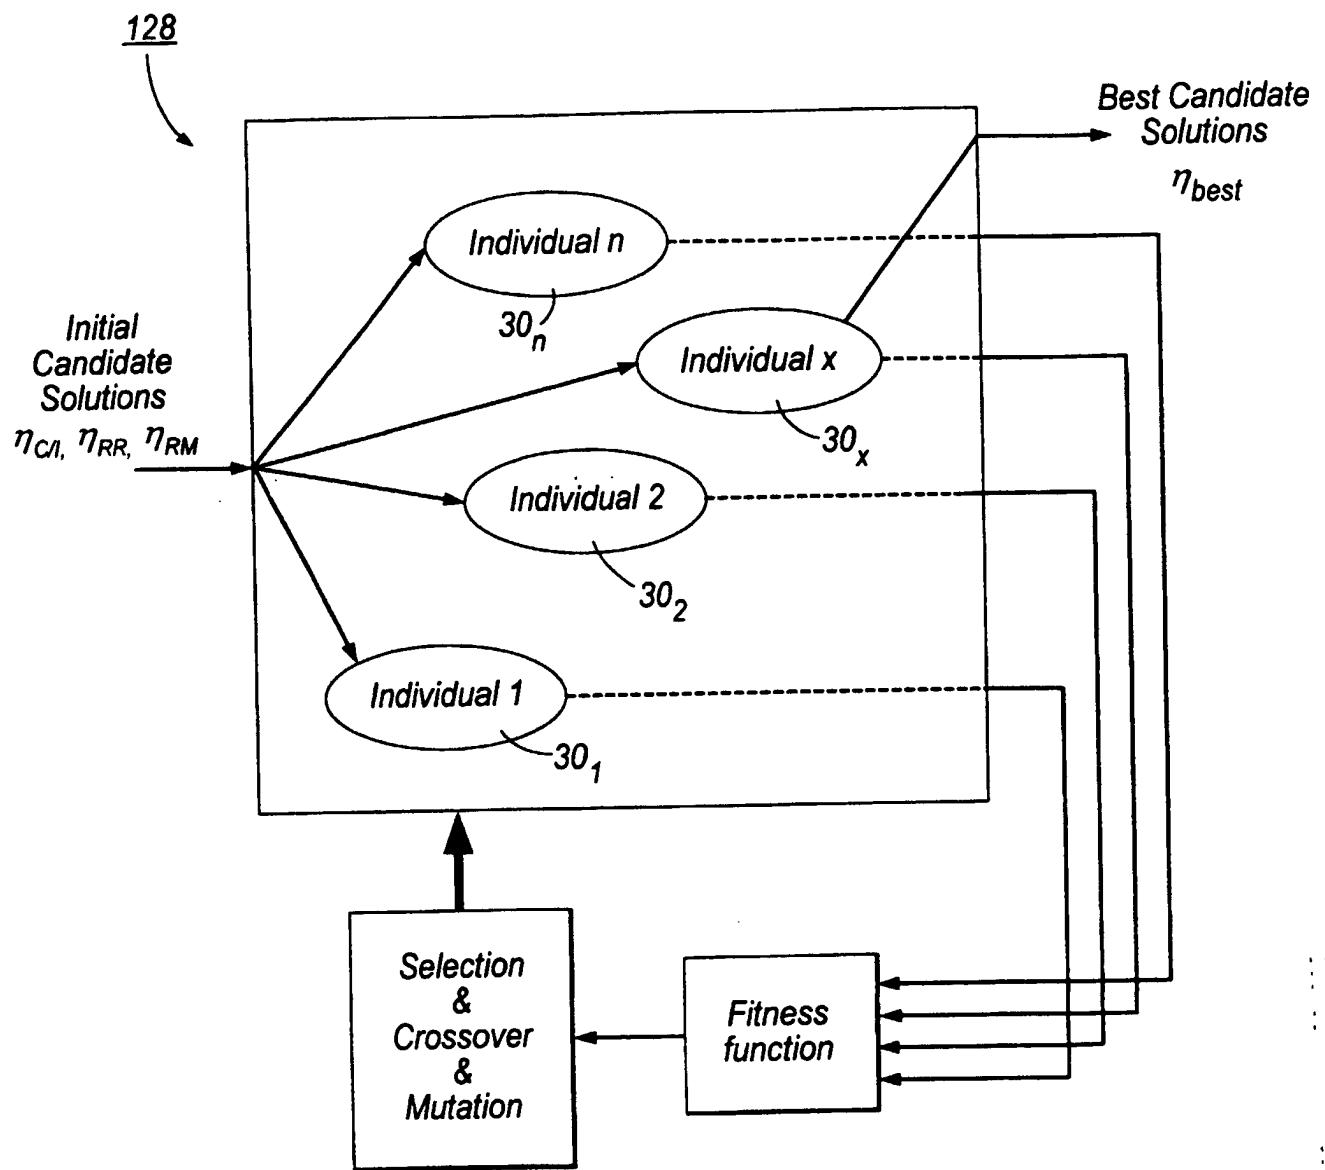


Fig.8

**Fig.9**

9/22

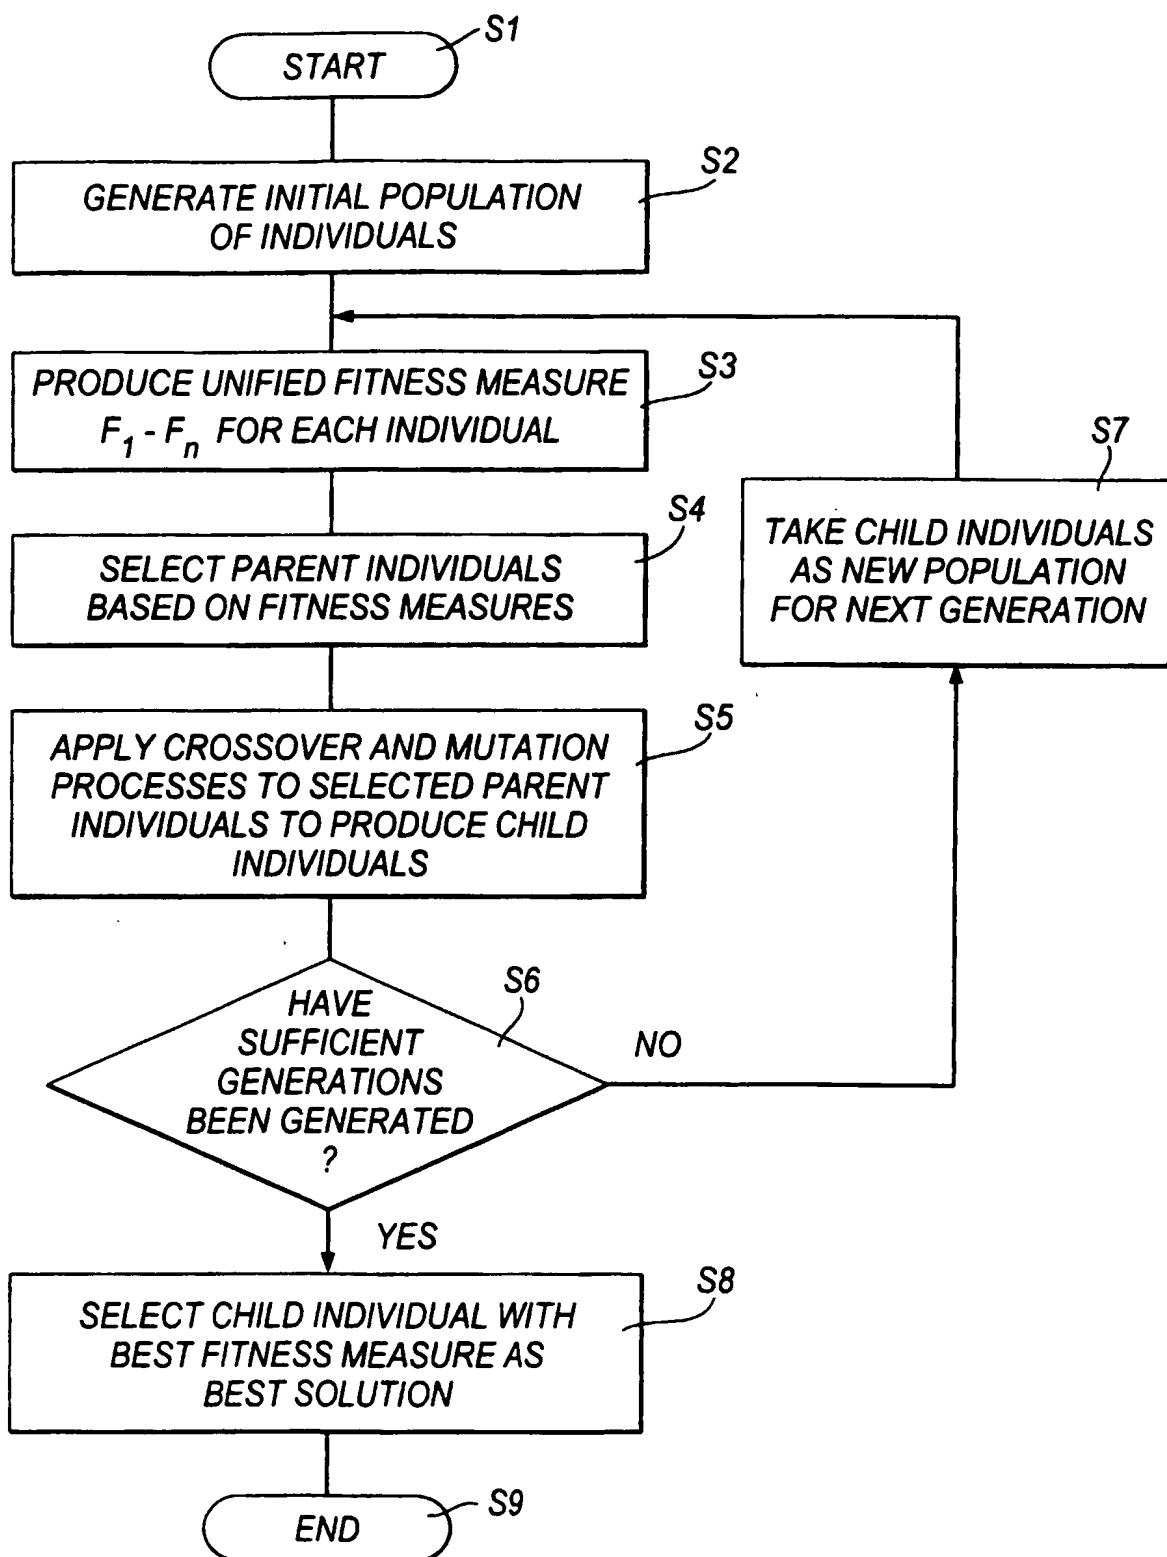


Fig.10

10/22

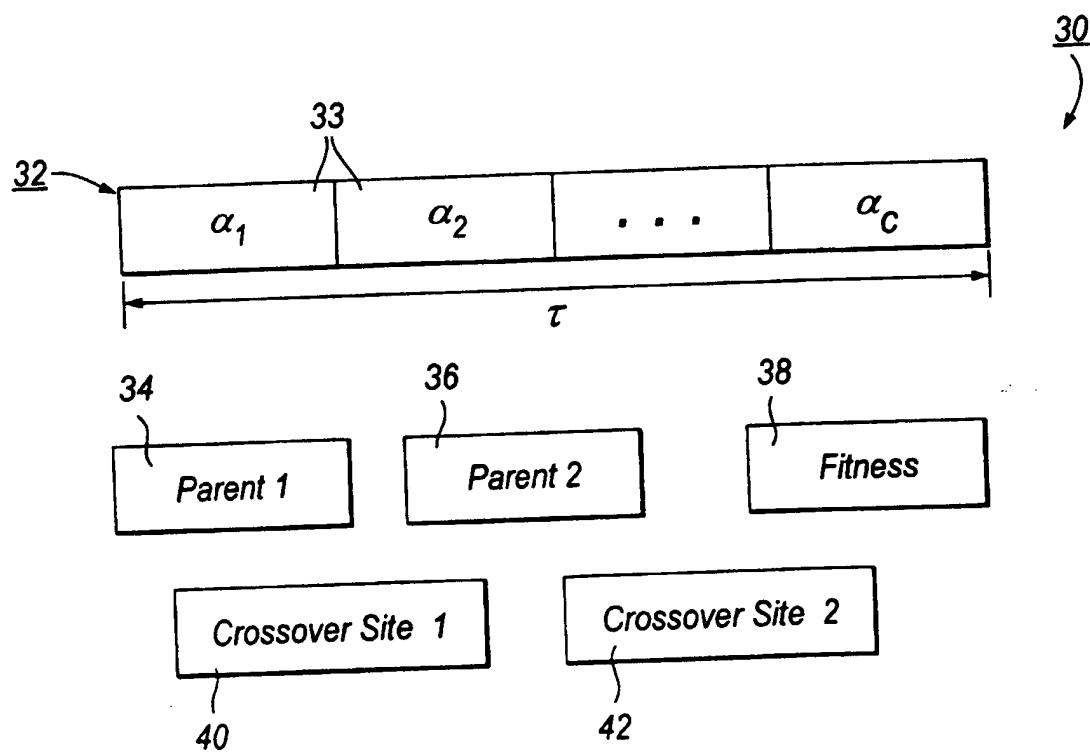
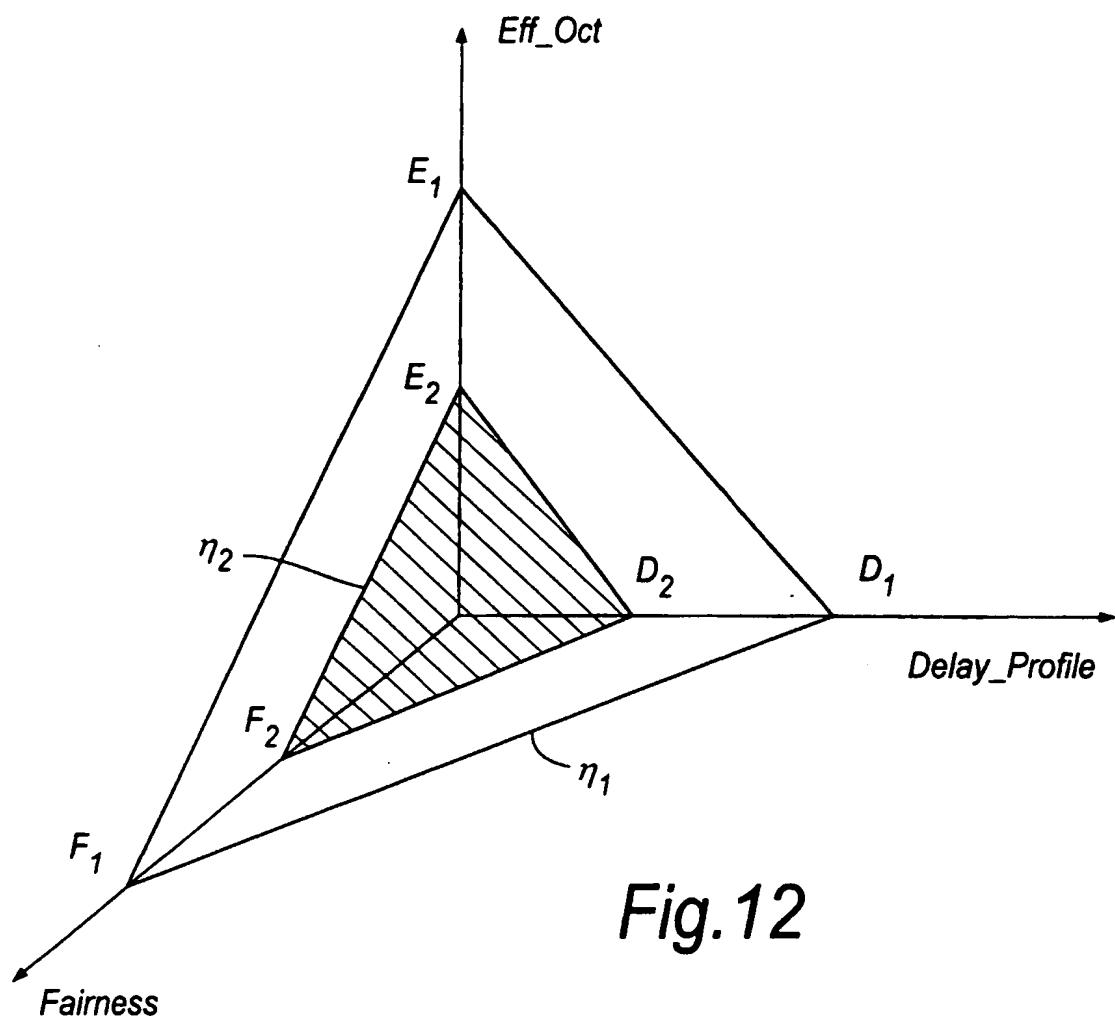


Fig. 11

11/22



(

12/22

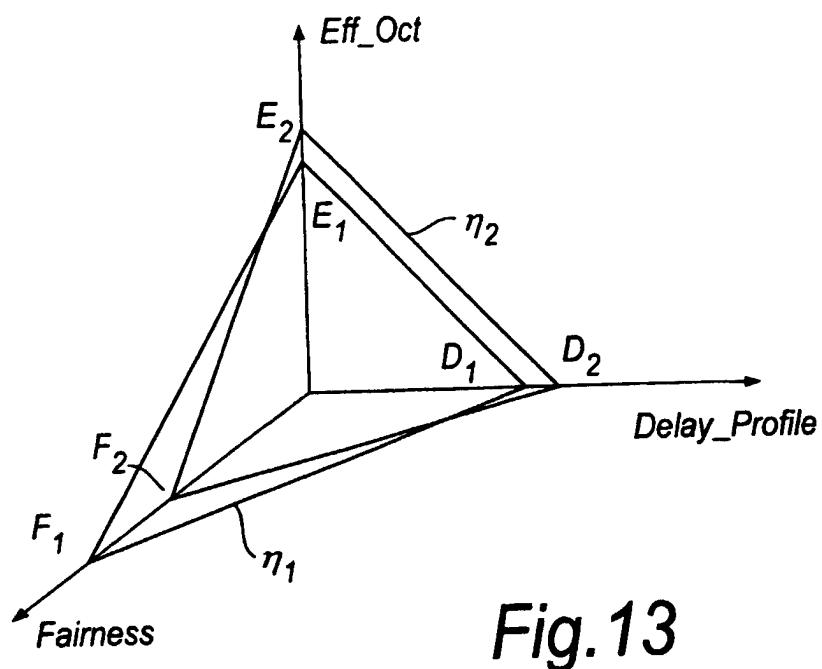


Fig.13

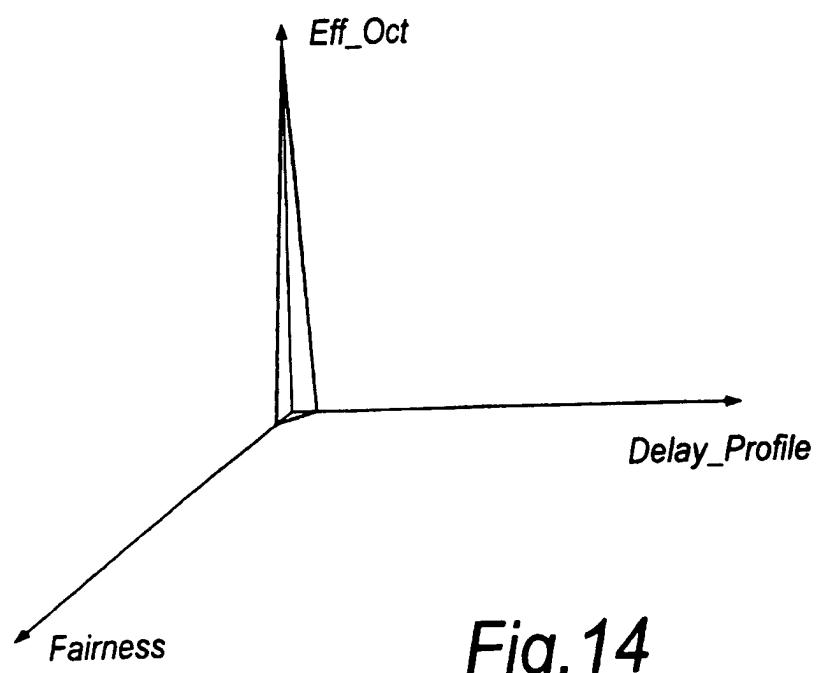


Fig.14

13/22

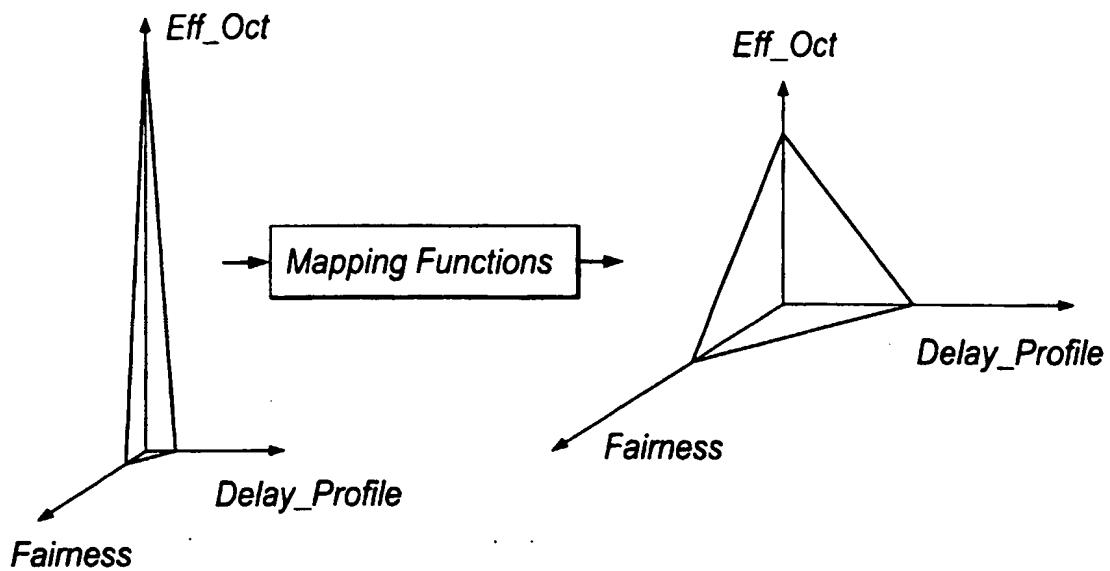


Fig.15

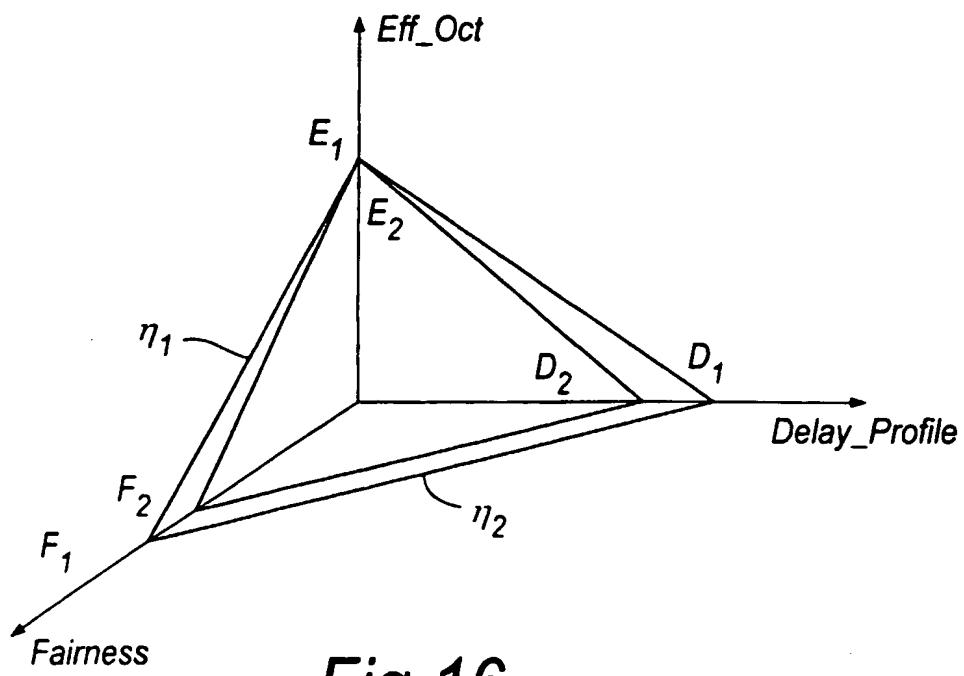


Fig.16

14/22

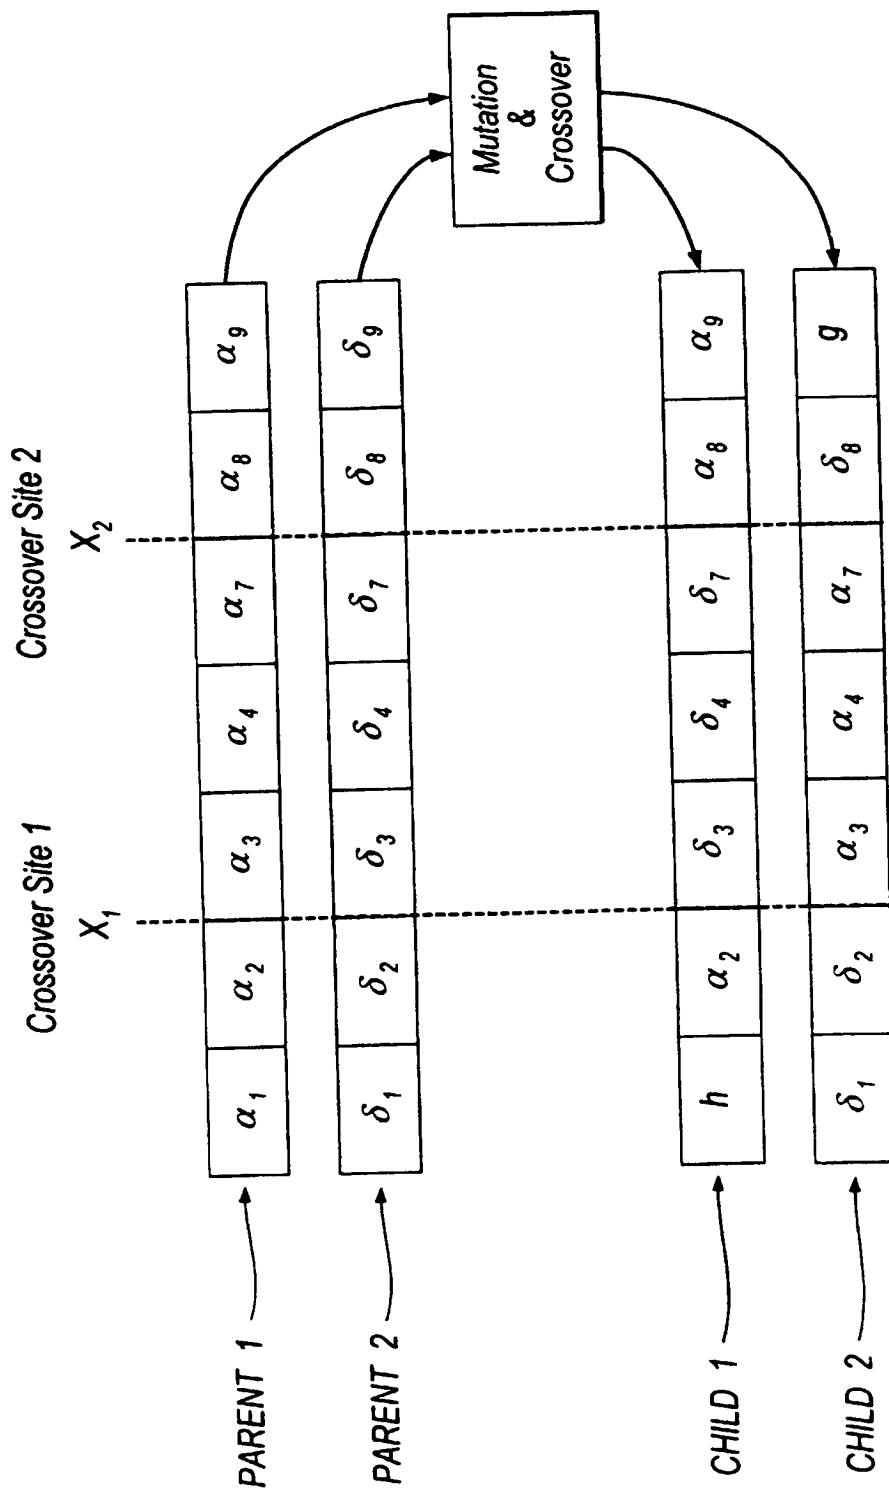


Fig. 17

15/22

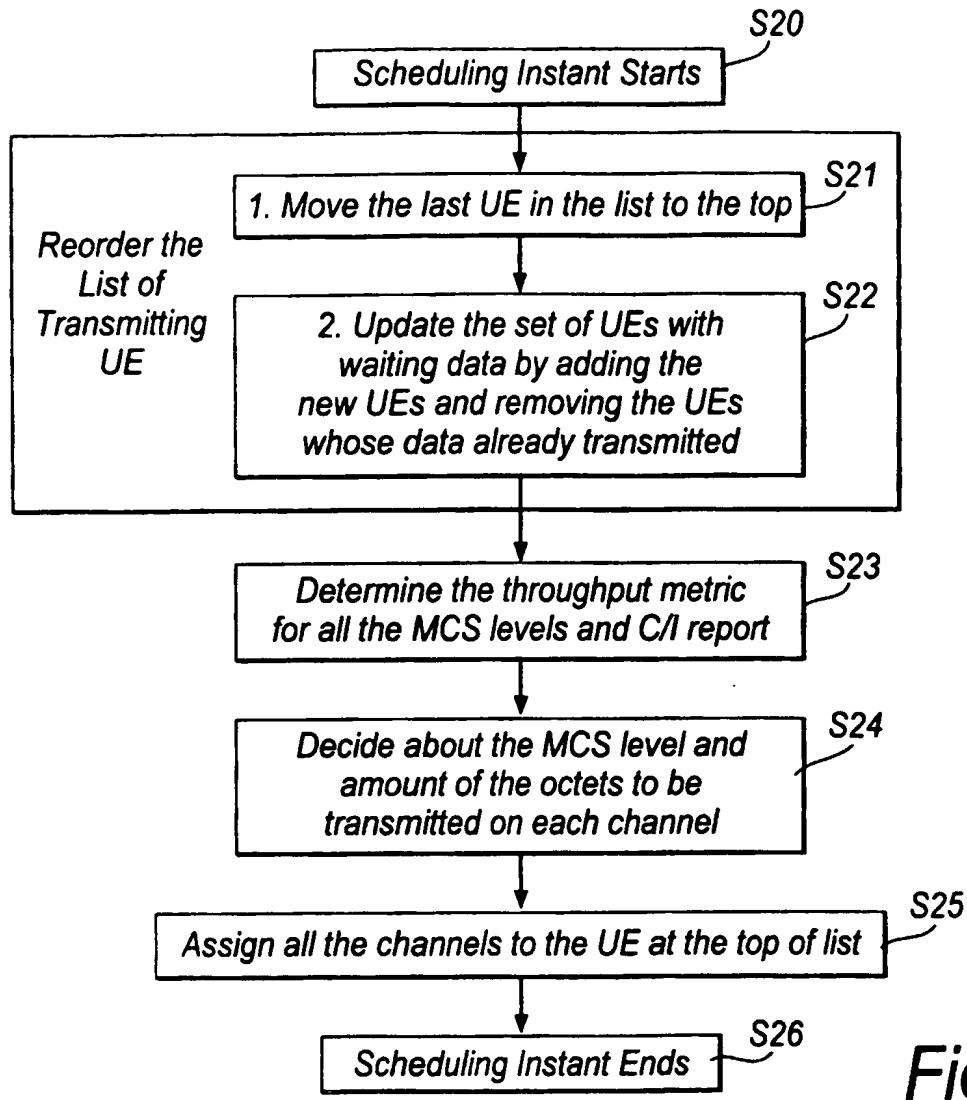


Fig. 18

UE 11
UE 5
UE 32
UE 1
UE 49
UE 8
UE 23
UE 34

Fig. 19(a)

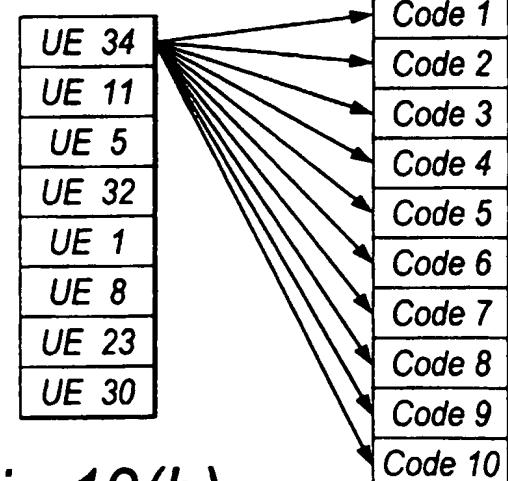


Fig. 19(b)

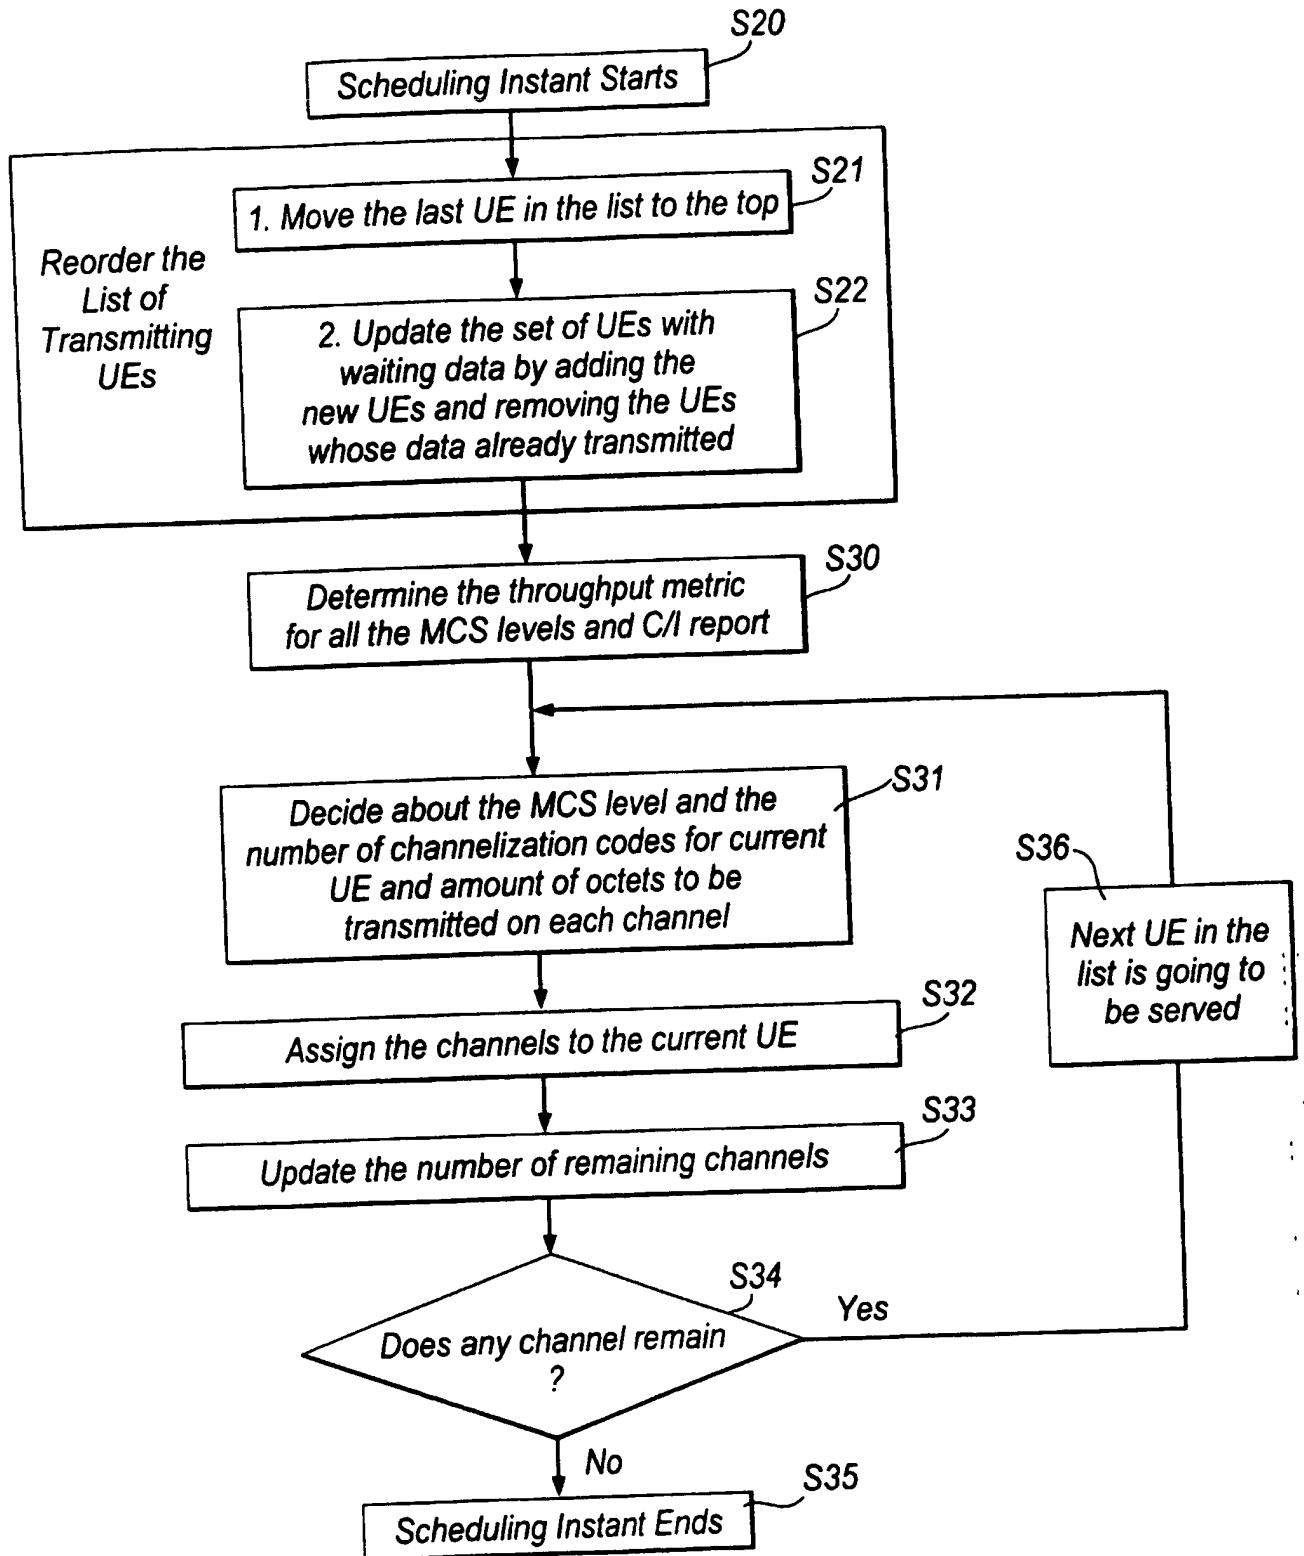


Fig.20

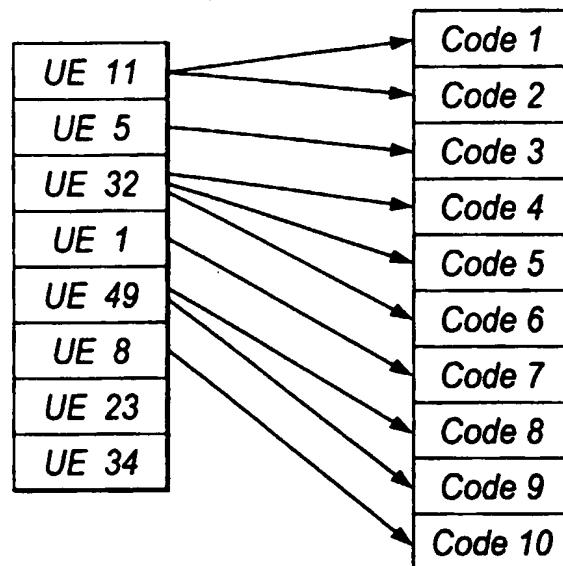


Fig.21(a)

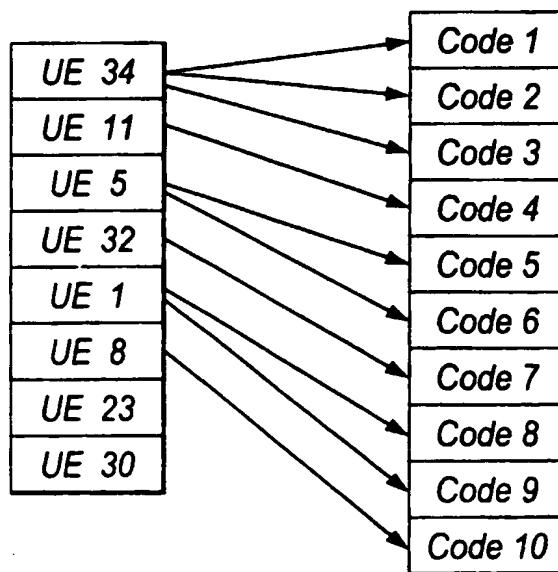


Fig.21(b)

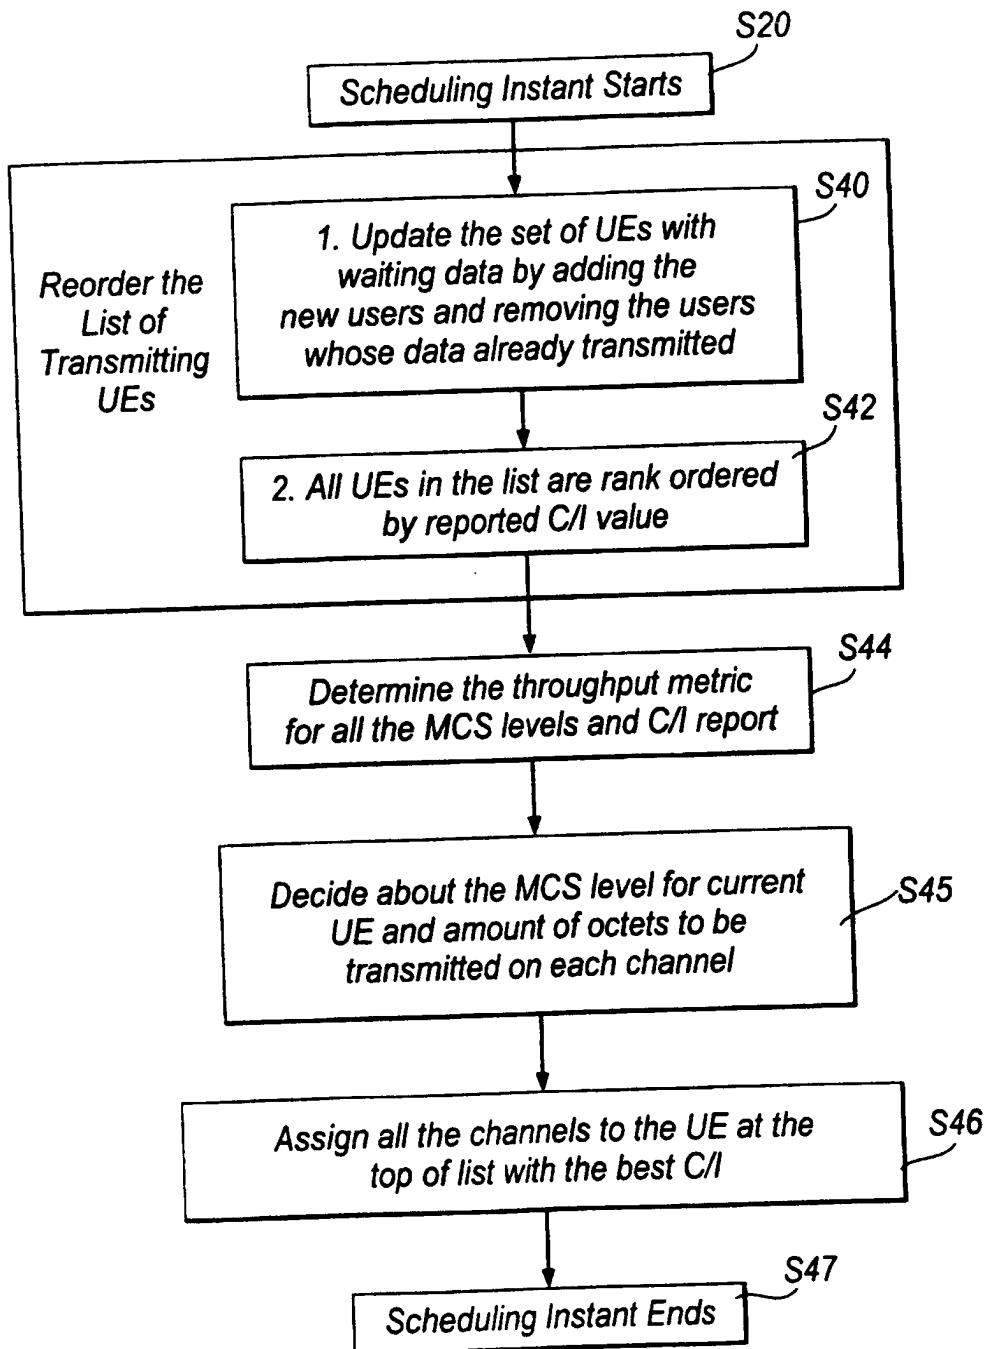


Fig.22

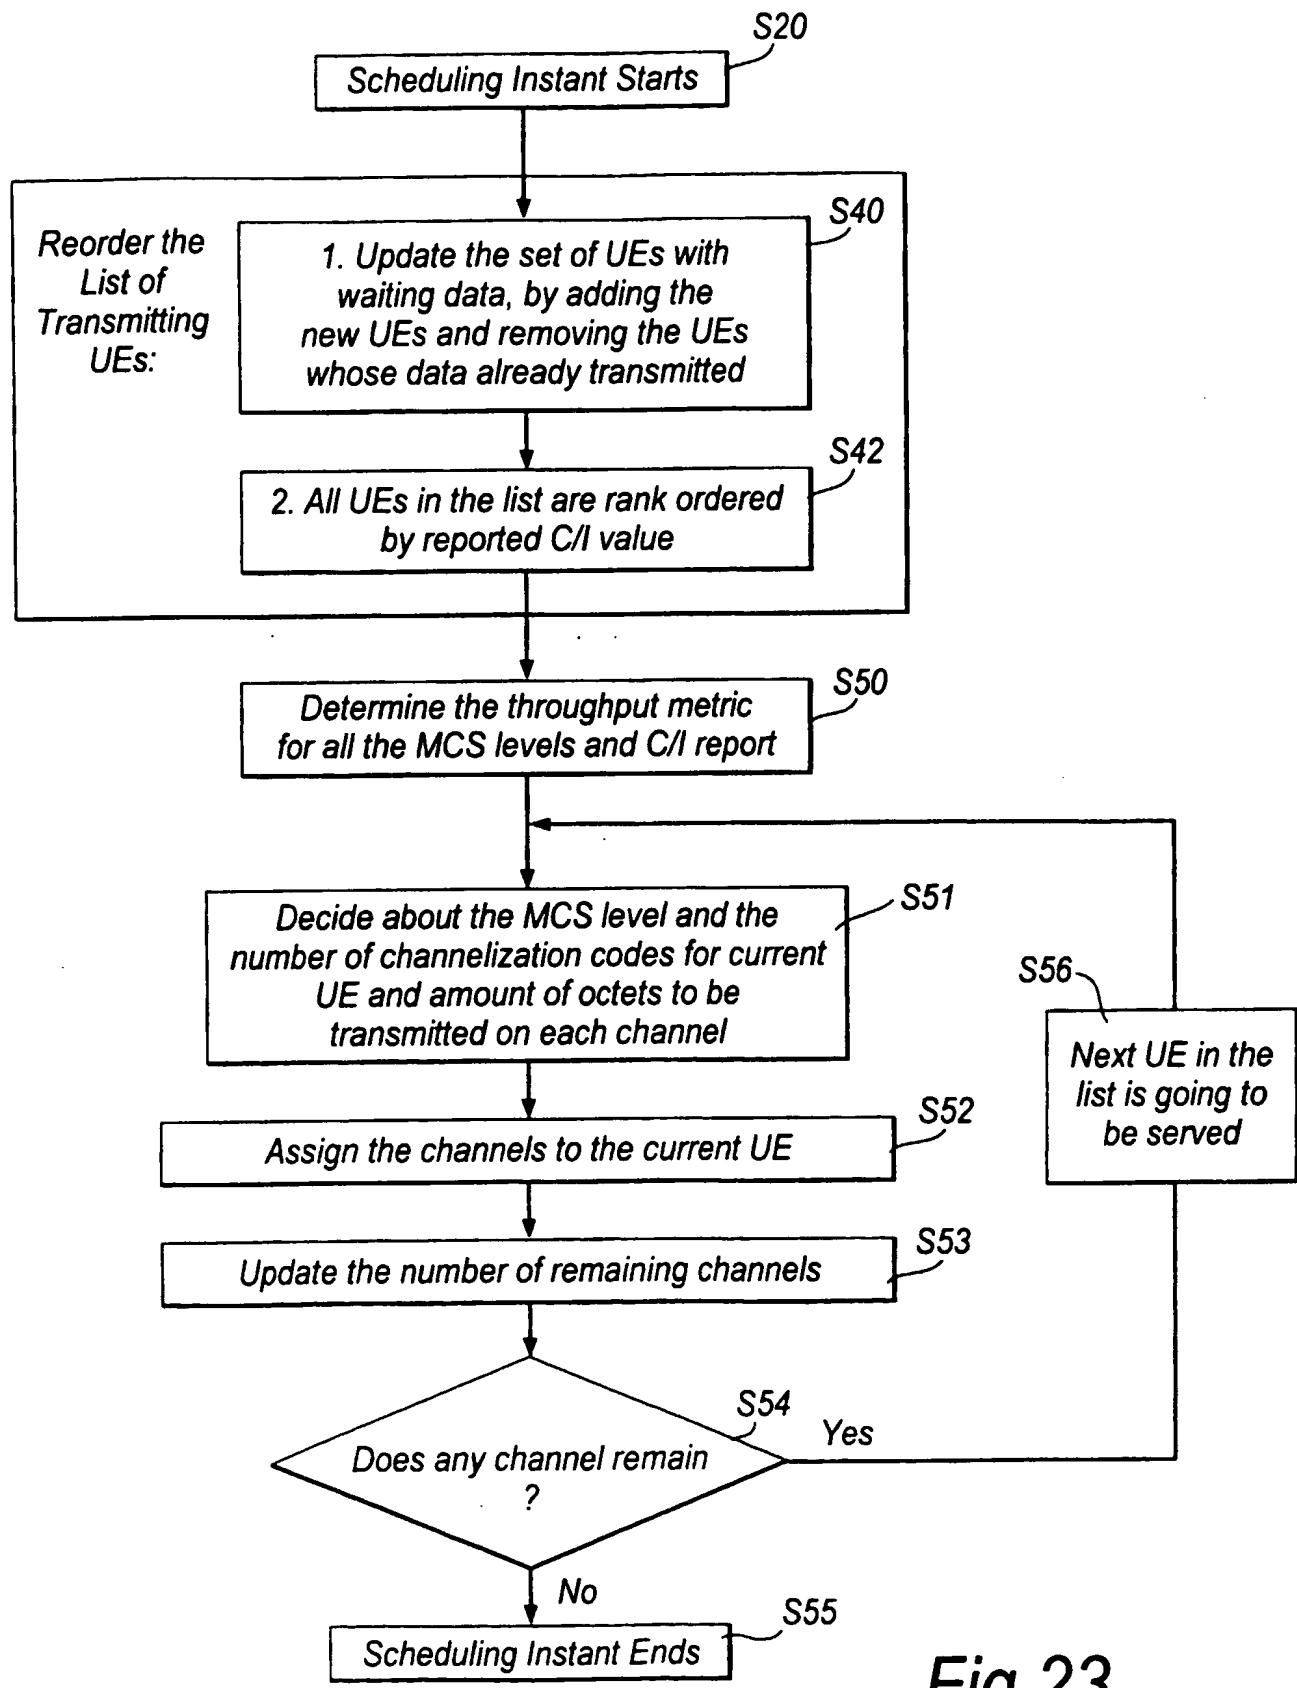


Fig.23

20/22

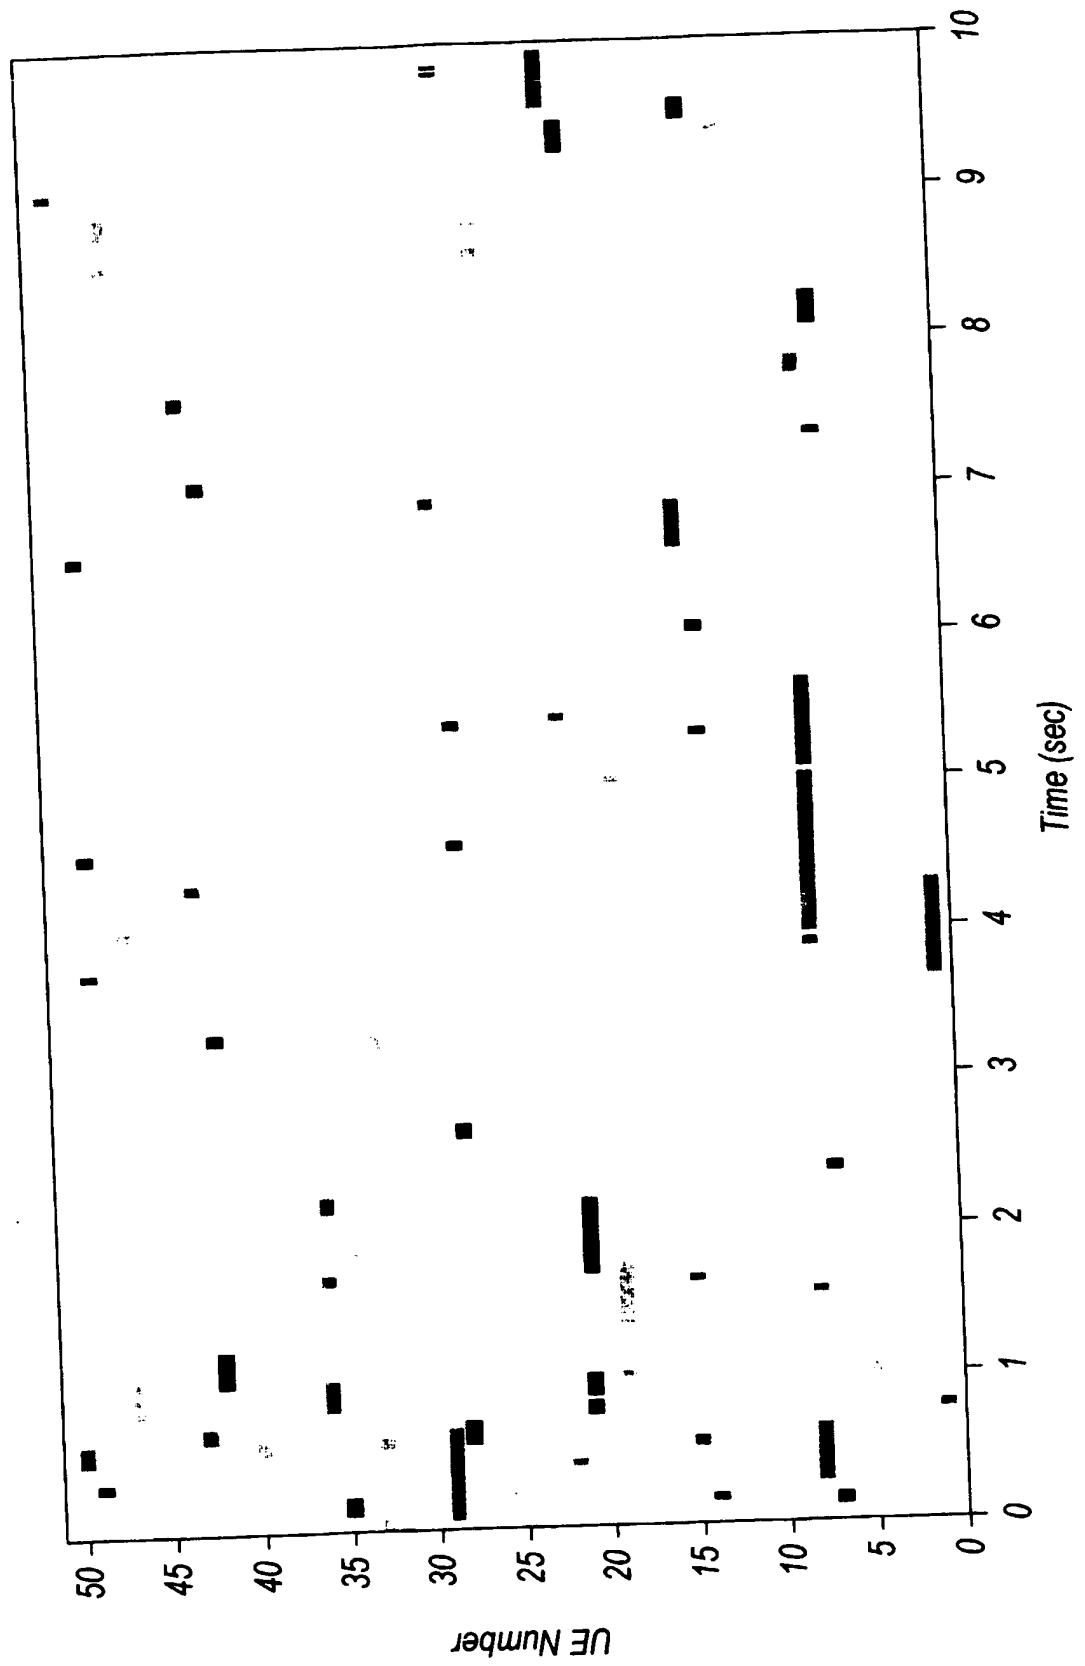


Fig.24

21/22

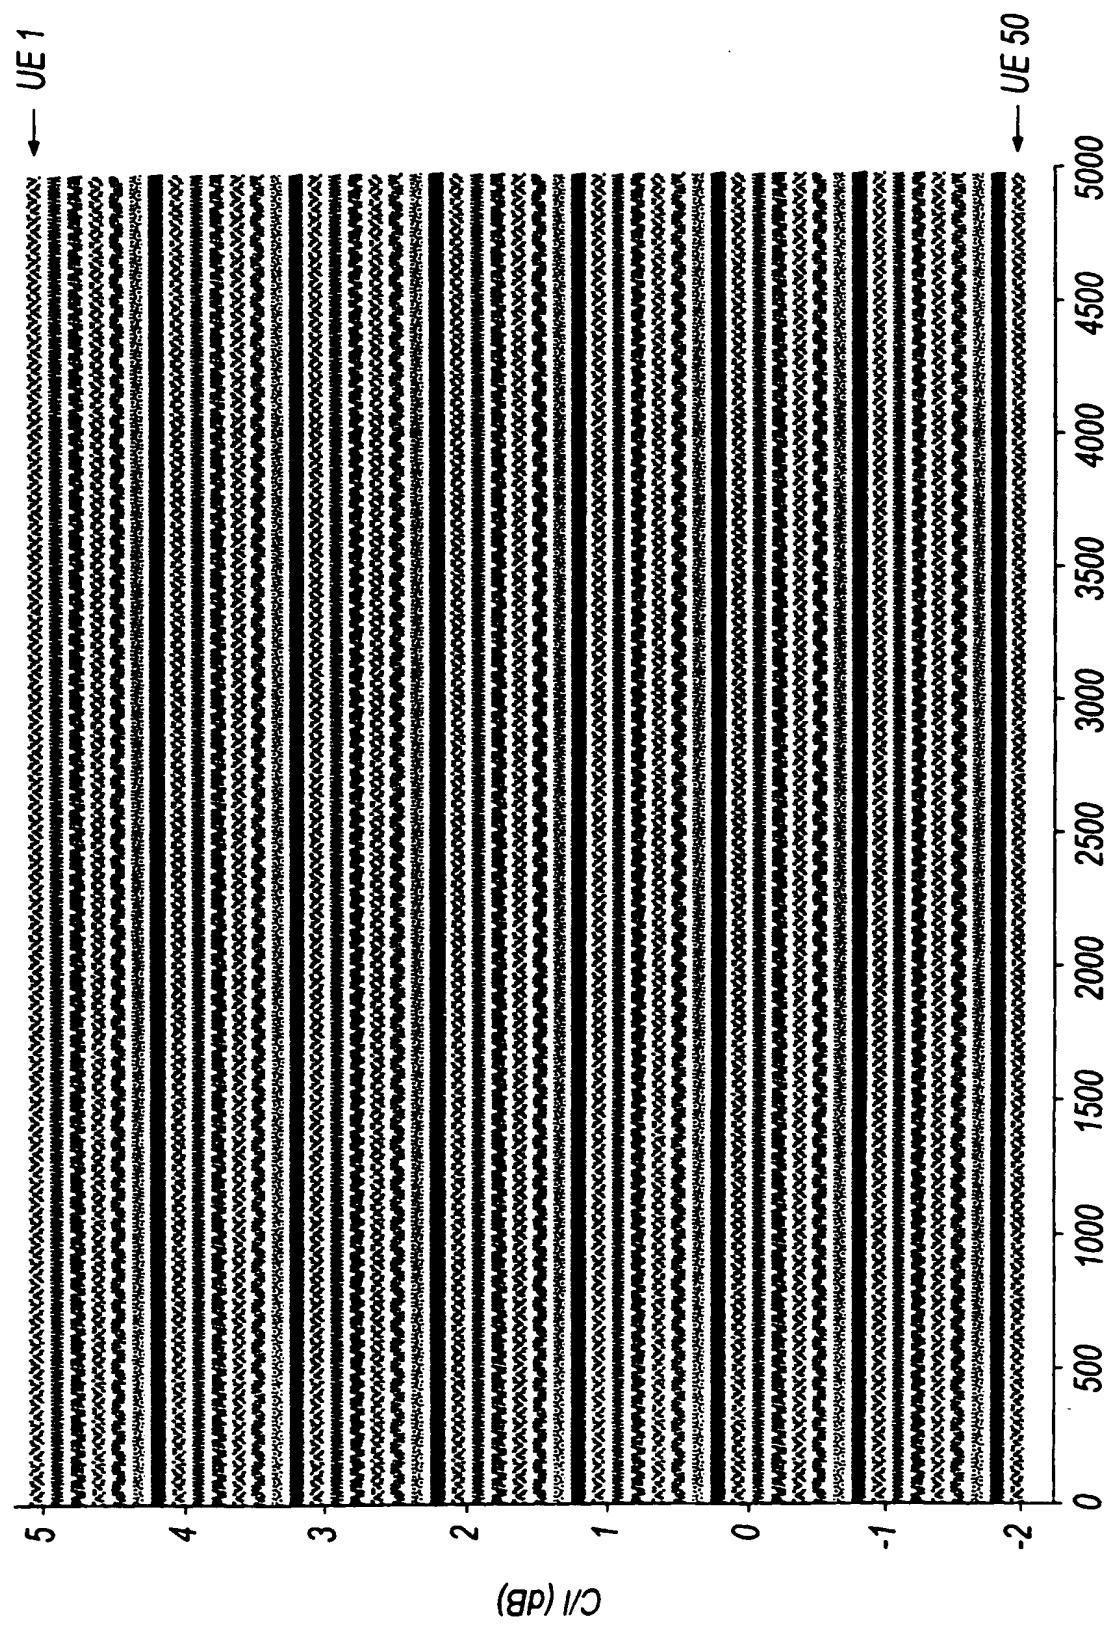


Fig. 25

TTI Number

22/22

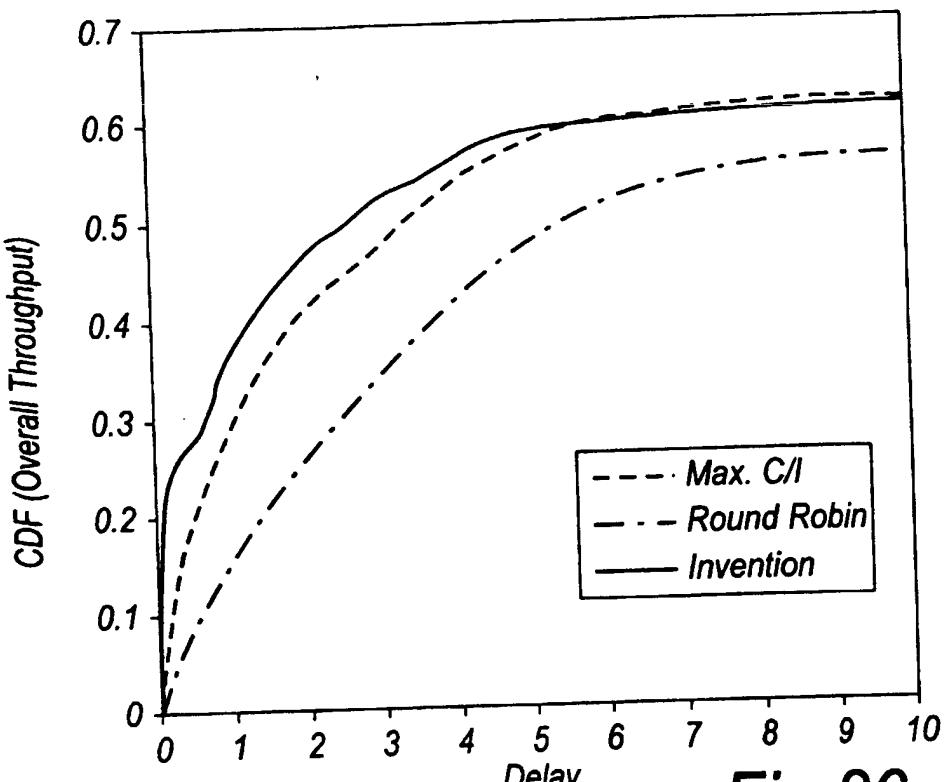


Fig.26

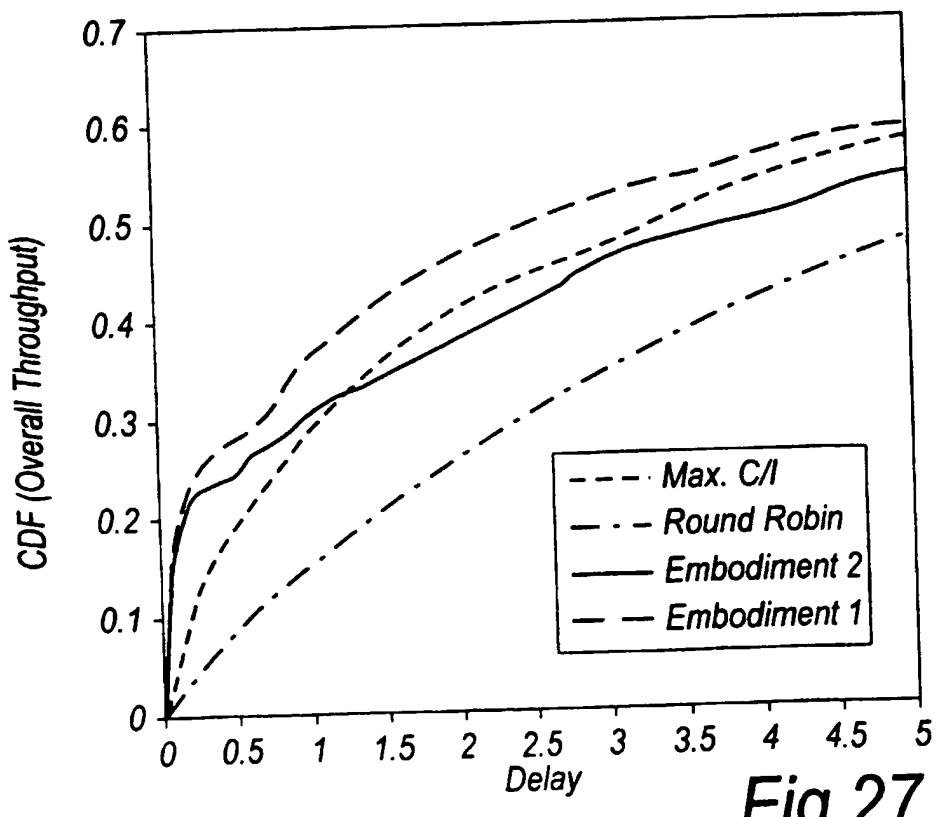


Fig.27

2390775

PACKET SCHEDULING

The present invention relates to packet scheduling methods and apparatus for use, for example, in wireless communication systems.

Fig. 1 shows parts of a wireless communication system 1. The system includes a plurality of base stations 2, only one of which is shown in Fig. 1. The base station 2 serves a cell in which a plurality of individual users may be located. Each user has an individual user equipment (UE). Only the user equipments UE1, UE2 and UE50 are shown in Fig. 1. Each UE is, for example, a portable terminal (handset) or portable computer.

As is well known, in a code-division multiple access (CDMA) system the signals transmitted to different UEs from the base station (also known as "node B") are distinguished by using different channelisation codes. In so-called third generation wireless communication systems a high speed downlink packet access (HSDPA) technique has been proposed for transmitting data in the downlink direction (from the base station to the UEs). In this technique a plurality of channels are available for transmitting the data. These channels have different channelisation codes. For example, there may be ten different channels C1 to C10 available for HSDPA in a given cell or sector of a cell. In HSDPA, downlink transmissions are divided up into a series of transmission time intervals (TTI), and a packet of data is transmitted on each different available channel to a selected UE. A new choice of which UE is served by which channel can be made in each TTI.

Fig. 2 shows an example of the operation of the HSDPA technique over a series of transmission time intervals TTI1 to TTI9. As shown in Fig. 2, in TTI1 it

is determined that two packets will be sent to UE50, four packets will be sent to UE11 and four packets will be sent to UE2. Accordingly, two channels are allocated to UE50 and four channels each are allocated to UE11 and UE2. Thus, as shown in Fig. 1, UE50 is allocated channels C1 and C2, UE11 is allocated channels C3 to C6, and UE2 is allocated channels C7 to C10.

In the next transmission time interval TTI2 a new user equipment UE1 is sent one packet, and the remaining UEs specified in TTI1 continue to receive packets.

Thus, effectively the HSDPA system employs a number of parallel shared channels to transmit data in packet form from the base station to the different UEs. This system is expected to be used, for example, to support world wide web (WWW) browsing.

In order to decide which UE should be served on which channel in each TTI a packet scheduling technique is employed. Conventionally, two basic types of scheduling technique have been considered for use in HSDPA: a round-robin (RR) scheduling technique and a maximum carrier-to-interference ratio (max C/I) technique.

The basic round-robin technique first compiles a list of the UEs which currently have data waiting at the transmitter (base station) for transmission. For each TTI the last UE in the list will have the highest priority for the next TTI. Accordingly, the UEs are serviced in a round robin fashion. In the simplest round-robin scheduling technique, it is assumed that the UE with the highest priority takes all of the channels. However, a packet-weighted round-robin technique is also known. This allocates the available channels to a group of users in each TTI based on the relative amounts of data for the different UEs. In

this packet-weighted technique UEs which have more data waiting for transmission are allocated more channels.

The round-robin scheduling techniques emphasise fairness amongst the competing UEs in terms of radio 5 resource allocation. However, they tend to provide relatively poor total throughput of data.

The max C/I scheduling technique is similar to the round-robin scheduling technique except that the list of UEs having waiting data is sorted in each TTI based 10 on a carrier-to-interference ratio (C/I) reported by each UE. The C/I is a measure of the quality of the channel. By sorting the list of UEs based on C/I, UEs which have a better channel quality are given a higher chance to be selected. In the simplest version of the 15 technique, all of the channels are allocated to the UE with waiting data that has the highest C/I. A packet-weighted variant is also possible, in which instead of selecting a single UE having the highest C/I, a group of UEs with the highest C/I values is selected, and the 20 available channels are divided up amongst the group of UEs based on the relative amounts of data which those UEs have waiting for transmission.

The max C/I scheduling technique tends to maximise the total throughput of data but this is at the expense 25 of fairness. It can be seen that UEs which report poor C/I values, for example because they are far from the base station or because there are many other interfering UEs in the vicinity, will only very rarely be selected. Thus, these UEs are likely to suffer from 30 unacceptably long delays in receiving packets.

Genetic algorithms have been considered for solving certain problems in the field of signal processing.

The principle of evolution is the primary unifying 35 concept of biology, linking every organism together in a historical chain of events. Every creature in the

(chain is the product of a series of "accidents" that have been sorted out thoroughly under selective pressure from the environment. While evolution has no intrinsic purpose (it is merely the effect of physical laws acting on and within populations and species), it is capable of engineering solutions to the problems of survival that are unique to each individual's circumstances. Taking a page from Darwin's "On the Origin of the Species", scientists have found ways to evolve solutions to complex problems.

10 The genetic algorithm was invented by John Holland in 1975 as one of the most powerful members of the class of stochastic search techniques, as described for example in "Genetic algorithms in search, optimisation, 15 and machine learning", D E Goldberg, Addison-Wesley, 1989. The mechanics of natural selection and genetics are used to create evolutionary optimisation. The best individuals from a generation are combined with a randomised exchange of information to create a stronger 20 next generation.

Each generation includes a population of individuals. Each of these individuals is defined by a string which is known as a chromosome. A string includes smaller units called genes.

25 For each individual, a unique fitness value is assigned. The parents of the next generation are selected based on the fitness values. Then these parents are mated under a process called crossover. To create new search paths every gene may have a random 30 change with a probability called the probability of mutation. Thus, the three major operators in a genetic algorithm are selection, crossover and mutation. In each generation, the fitness values become better and better.

35 One application of a genetic algorithm proposed by the present inventor is to suppress multiple access

(interference (MAI) and intersymbol interference (ISI) in code division multiple access (CDMA) communication systems, see for example "Genetically Modified Multiuser Detection for Code Division Multiple Access Systems", S Abedi & R Tafazolli, IEEE Journal on Selected Areas in Communications, vol. 20, no. 2, February 2002, pp. 463-473. In this application, termed Hybrid Genetic Multiuser Detection or HGMD, the problem of multiuser detection in a CDMA communication system is addressed using a genetic algorithm. In this application the genetic algorithm is embodied in a detector which detects the signal of a selected user and rejects interference from other users' signals. For example in a linear multiuser detector arrangement, a filter is defined by a number of taps, and transmitted information symbols in the selected user's signal are estimated based on the output of the filter. The genetic algorithm adjusts the tap coefficients of the filter over a series of iterations (generations).

Another application of a genetic algorithm proposed by the present inventor is a predictor for narrowband interference suppression in CDMA communication systems. This predictor, termed a Hybrid Genetic Predictor or HGP, is described in S Abedi, "Genetic Multiuser Detection for Code Division Multiple Access Systems, PhD Thesis, University of Surrey, October 2000.

Heretofore, there has been no practical proposal for applying a genetic algorithm to the problem of packet scheduling.

It is therefore desirable to provide a packet scheduling technique using a genetic algorithm which exploits the problem-solving ability of this class of algorithm effectively.

It is also desirable to provide a packet scheduling technique using a hybrid approach combining

a genetic algorithm with a conventional scheduling technique such as the round-robin or max C/I scheduling technique.

It is also desirable to provide a packet scheduling technique in which no one performance aspect (e.g. throughput or fairness) dominates above all others.

It is also desirable to provide a scheduling technique in which a balance between different 10 performance aspects (such as throughput, delay and fairness) can be selected by an operator of the system.

It is also desirable to provide a method of determining a unified measure of fitness of a candidate scheduling solution which can enable different 15 candidate scheduling solutions for a particular scheduling instant (transmission time interval) to be compared and fitter ones of the candidate solutions to be selected.

According to a first aspect of the present 20 invention there is provided a packet scheduling method, for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which method comprises: generating a plurality of candidate scheduling solutions, at least 25 one said candidate solution being generated using a genetic algorithm, and each candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration; and comparing the generated candidate 30 solutions and, based on the comparison results, selecting a best one of the candidate solutions to use to transmit packets in said scheduling instant.

Such a method can arrive at a superior scheduling 35 solution to conventional scheduling methods such as round-robin and max C/I using an acceptable level of computational resources.

Preferably the genetic algorithm operates over a series of iterations and has a plurality of individuals, each representing one of the plurality of candidate scheduling solutions, and in each iteration 5 parent individuals of a current generation produce child individuals of a next generation according to the genetic algorithm.

Each individual may have a chromosome comprising genes representing respectively parameters of the 10 candidate scheduling solution, and the genetic algorithm may process the genes of the parent individuals of the current generation to produce the genes of the child individuals of the next generation.

In a first generation at least one of the 15 individuals represents a candidate solution generated using a deterministic scheduling method. This enables the deterministic solution(s) to form part of the initial population and thereby influence subsequent genetic solutions.

20 The individuals in the first generation may include an individual representing a round-robin candidate scheduling solution generated by a round-robin scheduling method and/or an individual representing a maximum carrier-to-interference ratio 25 candidate solution generated by a maximum carrier-to-interference ratio scheduling method.

In a first generation at least one of the individuals may represent a randomly- or pseudorandomly-generated candidate solution. This can 30 increase the diversity of the initial population, which is an important element of the success of a genetic algorithm.

In many practical situations, for example in a wireless communication system, a plurality of channels 35 will be available for transmitting packets from the transmitter to the receivers in each scheduling

instant. In this case, each candidate solution preferably further specifies one or more parameters applicable to each available channel in the scheduling instant under consideration.

5 For example, each candidate solution may specify how the specified receiver(s) is (are) to be allocated to the available channels in the scheduling instant under consideration. In this case, each gene may correspond to one of the available channels and may 10 specify the receiver allocated to its corresponding channel in the scheduling instant under consideration.

Alternatively, or in addition, each candidate solution may specify a modulation and/or coding scheme (e.g. MCS level) to be applied to each available 15 channel in the scheduling instant under consideration. This is advantageous if the communication system uses an adaptive modulation scheme to increase the throughput when the channel conditions permit.

20 Alternatively, or in addition, each candidate solution may specify a transmission power for each available channel in the scheduling instant under consideration.

25 Alternatively, or in addition, each candidate solution may specify an amount of data to be transmitted via each available channel in the scheduling instant under consideration.

The comparison of the candidate solutions may be carried out in any suitable way. Preferably, however, for each candidate solution, at least one fitness 30 measure is determined, the or each fitness measure being a measure of performance of the candidate solution concerned.

35 Preferably, a set of two or more such individual fitness measures is determined for each candidate solution, which fitness measures relate to different respective performance aspects. This can enable the

method to balance two or more different performance aspects such as throughput and fairness. Conventional scheduling methods either explicitly or implicitly concentrate exclusively on one performance aspect and 5 ignore other aspects.

Preferably, the two or more individual fitness measures in the set for each candidate solution are combined together using a fitness function to produce a unified measure of fitness for the candidate solution 10 concerned. This enables a simple comparison of unified measures to be used to select the best solution, whilst preserving the ability to take into account a balance of performance aspects.

Preferably, the genetic algorithm is operable to 15 rank the individuals of a generation based on the respective unified fitness measures for the candidate solutions represented by these individuals, so that the ranking takes into account a balance of performance aspects.

20 The unified measure of fitness may be produced by forming a product of the individual fitness measures in the set. This is simple but effective. However, when forming the unified measure in this way it may be desirable to adopt one or more of the following 25 techniques to enhance the result.

Firstly, at least one the individual fitness measure of the set may be adjusted to reduce a mismatch between it and another one of the individual fitness measures of the set.

30 For example, for one or more the individual fitness measures of the set, a mapping function may be applied to the fitness measure so as to reduce a mismatch between it and another one of the individual fitness measures of the set.

Alternatively, or in addition, one or more of the individual fitness measures of the set may be normalised.

Secondly, one or more of the individual fitness measures of the set may be weighted when the individual fitness measures are combined together to produce the unified measure of fitness. This enables a desired balance between different performance aspects to be achieved. The weights may be adjustable, if required, to provide a tuning facility and enable operators to choose individually-preferred balances.

The performance aspects which can be taken into account are not limited. However, the unified fitness measure may be based on one or more of the following:

a measure of total data throughput to the specified receiver(s) for the candidate solution concerned;

a measure of an amount of data at the transmitter waiting for transmission to the specified receiver(s) for the candidate solution concerned;

a measure of delay(s) in transmitting data to the specified receiver(s) in the candidate solution concerned; and

a measure of how the candidate solution affects a fairness of the packet scheduling amongst active ones of the plurality of receivers.

To cope with different service types (e.g. video data, WWW download data etc.) without undue complexity, the or each delay is preferably calculated relative to a preselected delay parameter applicable to a type of data to be transmitted to the receiver concerned. For example, the delay parameter applicable to a higher-priority type of data may be lower than the delay parameter applicable to a lower-priority type of data.

A measure of total throughput may be calculated by estimating a total amount of data which will be

transmitted successfully to the specified receiver(s) of the candidate solution in the scheduling instant under consideration.

A measure of waiting data may be a ratio of a current amount of data waiting for transmission at the transmitter to each specified receiver to a total the amount of data which has arrived at the transmitter for transmission to the receiver concerned.

A measure of delay may be calculated based on a sum of respective delay times for transmissions to the specified receiver(s). Each delay time may be dependent upon a difference between a current time and a time of arrival at the transmitter of the earliest data currently waiting for transmission to the specified receiver.

A fairness measure may be based on a variance or average of respective estimated data throughputs to all active ones of the receivers if the candidate solution is used.

A measure of quality of service may be based on the number of packets delivered to each receiver within a desired threshold of time. The threshold may be different for different service types.

The type of transmitter and receiver in the invention is not limited but in one embodiment the transmission is a wireless transmission and the transmitter is part of a base station of a wireless communication system, and each receiver is part of a user equipment of that system. A "packet" may be of any suitable length.

The genetic algorithm may adopt sexual or asexual reproduction models from nature, but is not constrained to use either of these models. In one embodiment the individuals of a generation are selected as parent individuals in dependence upon their rankings.

(In an embodiment based on the sexual reproduction model the genetic algorithm causes two parent individuals to have two child individuals and the genes of the child individuals are dependent on the genes of their parent individuals.

5 The parent-individual chromosomes may be subject to crossover and/or mutation processes to generate the child-individual chromosomes.

Preferably, in the crossover process at least one 10 subset of the genes of each chromosome is defined and at least one gene in the subset in the first child individual is derived from a corresponding gene in the subset in the second parent individual and at least one gene in the subset in the second child individual is 15 derived from a corresponding gene in the subset in the first parent individual.

Preferably, at least one gene not in the subset in the first child individual is derived from a corresponding gene not in the subset in the first 20 parent individual, and at least one gene not in the subset in the second child individual is derived from a corresponding gene not in the subset in the second parent individual.

The or each such subset, if any, may be selected 25 randomly or pseudorandomly.

In the mutation process at least one gene in at least one child individual may be subject to a change. The gene may be changed to specify a receiver having data waiting at the transmitter.

30 According to a second aspect of the present invention there is provided packet scheduling apparatus, for scheduling packets of data for transmission from a transmitter to a plurality of receivers via at least one channel, which apparatus 35 comprises: solution generating means for generating a plurality of candidate scheduling solutions, at least

one said candidate solution being generated using a genetic algorithm, and each said candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration; and best solution selecting means for comparing the generated candidate solutions and selecting, based on the comparison results, a best one of the candidate solutions to use for transmitting packets in the scheduling instant.

According to a third aspect of the present invention there is provided a transmitter comprising: packet scheduling apparatus embodying the second aspect of the invention; and transmitting means connected operatively to said packet scheduling apparatus and operable, when said scheduling instant occurs, to cause packets to be transmitted to the receiver(s) specified by the best solution selected for that instant by the packet scheduling apparatus.

According to a fourth aspect of the present invention there is provided a packet scheduling method, for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which method comprises: generating a plurality of candidate scheduling solutions, each candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration; producing a measure of fitness for each said candidate solution, said fitness measure being a measure of performance of the candidate solution in at least one of the following aspects: delay, fairness, and quality of service; and employing the measures of fitness produced for different candidate solutions to make a selection amongst the generated candidate solutions.

According to a fifth aspect of the present invention there is provided a packet scheduling method,

(for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which method comprises: generating a plurality of candidate scheduling solutions, each 5 candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration; producing a measure of fitness for each said candidate solution, said fitness measure being a measure of performance of the candidate 10 solution in at least two different aspects; and employing the respective measures of fitness produced for different candidate solutions to make a selection amongst the generated candidate solutions. The different performance aspects may include two or more 15 of the following: throughput, delay, fairness, and quality of service.

Preferably, in producing the fitness measure a weighting of at least one the performance aspect is adjustable relative to that of another the performance 20 aspect.

Each performance aspect preferably has an individually-adjustable weighting.

The fitness measures are preferably employed to select a best one of the candidate solutions to use to transmit packets to the receivers in the scheduling 25 instant under consideration.

At least one of the candidate solutions may be generated using a genetic algorithm and the fitness measures are employed in a selection process of the 30 genetic algorithm. Alternatively, some or all of the candidate solutions may be generated in other ways, for example based on random or deterministic solutions as described in our co-pending United Kingdom patent application no. [agent's ref. HL823171],^{G80216239.4} the entire 35 content of which is incorporated herein by reference.

Reference will now be made, by way of example, to the accompanying drawings, in which:

Fig. 1 discussed hereinbefore, shows parts of a wireless communication system employing a HSDPA technique for downlink transmissions;

Fig. 2 shows an example of the operation of the HSDPA technique in the Fig. 1 system;

Fig. 3 shows a block diagram of packet scheduling apparatus embodying the present invention;

Fig. 4 is a block diagram showing an example of the constitution of one of the parts of the Fig. 3 apparatus;

Fig. 5 is a block diagram showing an example of the constitution of another of the parts of the Fig. 3 apparatus;

Fig. 6 is a graph illustrating a example variation in a carrier-to-interference ratio of a downlink channel over a series of transmission time intervals for four different UEs in a wireless communication system;

Fig. 7 is a graph for use in explaining an adaptive modulation and coding technique suitable for use in embodiments of the invention;

Fig. 8 is a schematic diagram for use in explaining an automatic repeat request process suitable for use in embodiments of the present invention;

Fig. 9 is a schematic view of a genetic algorithm unit suitable for use in an embodiment of the present invention;

Fig. 10 is a flowchart illustrating operation of the Fig. 9 embodiment;

Fig. 11 is a schematic view for use in explaining a structure of an individual in the Fig. 9 embodiment;

Figs. 12 to 16 are schematic views for use in explaining a unified fitness measure used in the Fig. 9 embodiment;

(Fig. 17 is a schematic view for use in explaining crossover and mutation processes performed in the Fig. 9 embodiment;

5 Fig. 18 is a flowchart for use in explaining a scheduling process performed by a generator part shown in Fig. 4;

Figs. 19(a) and 19(b) are schematic views showing an example of the Fig. 18 process;

10 Fig. 20 is a flowchart for use in explaining a variation of the Fig. 18 process;

Figs. 21(a) and 21(b) are schematic diagrams showing an example of the operation of the Fig. 20 process;

15 Fig. 22 is a flowchart for use in explaining a scheduling process performed by another generator part shown in Fig. 4;

Fig. 23 is a flowchart for use in explaining a variation of the Fig. 22 process;

20 Fig. 24 is a diagram illustrating an example of packet transmission activity in an embodiment of the present invention;

Fig. 25 is a schematic diagram for use in explaining carrier-to-interference ratio differences applied to different UEs in a simulation process;

25 Fig. 26 is a graph illustrating a cumulative density function of overall throughput of data versus delay characteristic for an embodiment of the present invention and for two conventional schedulers; and

Fig. 27 is a larger-scale version of a portion of
30 the Fig. 26 graph.

Fig. 3 shows a block diagram of packet scheduling apparatus 10 embodying the present invention. The apparatus 10 is used to schedule packets of data for transmission from a transmitter to a plurality of receivers via at least one channel. The transmitter is, for example, a base station (node B) in a wireless

communication system. The plurality of receivers in this case are different UEs served by the base station.

The apparatus 10 comprises a candidate solution generator unit 12 which generates a plurality (η) of candidate scheduling solutions η_1 to η_n . Each candidate solution η specifies at least the receiver(s) to which packets are to be transmitted in a scheduling instant (e.g. transmission time interval TTI) under consideration. At least one of the candidate solutions is generated using a genetic algorithm.

The apparatus 10 also comprises a best solution selecting unit 14 which compares the different candidate solutions $\eta_1 \dots \eta_n$ generated by the generating unit 12 and selects a best one of the candidate solutions η_{best} based on the comparison results.

The apparatus 10 is connected operatively to a transmission unit 16 which, when a scheduling instant for which the apparatus has selected a best solution occurs, causes packets to be transmitted to the receiver(s) specified in the selected best solution η_{best} .

In the usual case, in which a plurality of channels are available for transmitting packets from the transmitter to the receiver(s), each candidate solution η must further specify how the specified receiver(s) is (are) to be allocated to the available channels in the scheduling instant under consideration.

Fig. 4 shows an example of the constitution of the candidate solution generator unit 12 in the Fig. 3 apparatus. In this example, the candidate solution generator unit 12 comprises first, second and third solution generators 122, 124 and 126, and a genetic algorithm unit 128.

The first generator 122 is a max C/I solution generator 122 which receives carrier-to-interference ratio (C/I) reports from each active receiver and

generates a max C/I candidate solution $\eta_{C/I}$ for the scheduling instant under consideration based on the reported C/I values. Details of the way in which the max C/I solution is generated by the generator 122 will be given later with reference to Figs. 18 to 21.

The second generator 124 is a round-robin (RR) solution generator which generates a candidate solution η_{RR} for the scheduling instant under consideration using a round robin scheduling technique. Details of the way in which the RR solution η_{RR} can be generated by the generator 124 will be given later with reference to Fig. 22 and 23.

Incidentally, the data destined for each different receiver is buffered in the transmitter prior to transmission, for example in a transmit queue or source queue corresponding to the receiver concerned. Depending on the type of scheduling technique desired, information FILL regarding the levels of fullness of the source queues for the different receivers may be provided to the first and second generators 122 and 124 so that the solutions $\eta_{C/I}$ and η_{RR} can take account of these fill levels. This is required, for example, when the solutions are packet-weighted, as described later with reference to Figs. 20 and 22.

The third generator 126 is a random solution generator which generates a random candidate solution η_{RM} in which the specified receiver(s) is (are) selected randomly or pseudo-randomly. For example, if there are N receivers the third generator 126 generates an integer between 1 and N for each available channel.

The genetic algorithm unit 128 generates one or more genetic candidate solutions by evolving a plurality (or population) of individuals over a series of iterations (or generations). Each individual represents one candidate solution. Parent individuals of one generation produce child individuals of a next

generation. After the individuals have evolved over a sufficient number of generations the genetic algorithm unit 128 outputs at least one genetic candidate solution as represented by an individual of the last 5 generation. The or each such genetic candidate solution is included in the plurality of candidate solutions η_1 to η_n generated by the candidate solution generating unit 12.

The genetic algorithm unit 128 requires an initial 10 (or seed) population of individuals to be defined. The initial population may be defined in many different ways. For example, as represented in Fig. 4, a first individual may represent the max C/I solution $\eta_{C/I}$ supplied by the first generator 122. A second 15 individual may represent the round-robin solution η_{RR} supplied by the second generator 124. Each further individual may represent a different random solution η_{RM} supplied by the third generator 126.

It is not necessary to define the initial 20 population in this way. For example, all the initial population could be random solutions. Alternatively, all the initial population could be derived from one or both of the max C/I solution and the round robin solution, for example by applying random or even 25 predetermined offsets to one or both of those solutions. In general, any initial population with individuals representing diverse solutions can be used.

The genetic solutions may make up the entire 30 plurality of candidate solutions η_1 to η_n output by the candidate solution generating unit 12. However, it is also possible to include one or both of the max C/I and round robin solutions in the plurality of candidate solutions η_1 to η_n together with the genetic solutions, so that these are available for comparison with the 35 genetic solutions. This could be worthwhile if, for

example, the initial population does not include one or both of the solutions $\eta_{C/I}$ and η_{RR} .

5 A detailed explanation of the operation of the genetic algorithm unit 128 in one embodiment of the invention will be given later with reference to Figures 9 to 17.

Fig. 5 shows an example of the constitution of the solution selecting unit 14 in the Fig. 3 apparatus. 10 Each candidate solution η_1 to η_n is supplied to a fitness measure producing unit 142 and a selector 146.

15 For each candidate solution η_1 to η_n the fitness measure producing unit 142 produces at least one fitness measure F_1 to F_n . The or each fitness measure for a candidate solution is a measure of performance expected of the candidate solution if it is actually chosen for use in the scheduling instant under consideration. As described later in more detail, preferably the fitness measure or measures should represent the expected performance of the candidate 20 solution in two or more difference performance aspects which it is desired to balance in the scheduling process. For example, the fitness measures may take into account one or more of the following performance aspects: throughput (average and total), delay and 25 fairness of packet delivery.

The fitness measures F_1 to F_n for the different candidate solutions η_1 to η_n are compared by a fitness measure comparing unit 144 which determines which one of the candidate solutions η_1 to η_n has the best fitness 30 measure (or best set of fitness measures) and outputs a control signal "best" indicating the identified solution. The selector 146 receives the control signal "best" and selects the solution η_{best} from among the candidate solutions η_1 to η_n .

35 Prior to describing preferred embodiments of the present invention in more detail, some background

information regarding the HSDPA system will be provided.

In the HSDPA system, channel state information (CSI) is made available to both the transmitter and the receiver, in order to realise a robust communication system structure. The HSDPA system is intended to increase the transmission rates and throughput, and to enhance the quality of service (QoS) experienced by different users. It transfers most of the functions from the base station controller (also known as the radio network controller or RNC) to the base transceiver station (node B).

In addition to employing a number of parallel shared channels as described in the introduction with reference to Fig. 2, the HSDPA system may also seek to achieve high transfer rates using other control techniques, as will now be explained.

One such control technique, referred to as a modulation and coding scheme (MCS), enables the transmitter to select different modulation and/or coding schemes under different channel conditions.

The C/I value for a channel between the transmitter and a receiver (UE) varies significantly over time. Fig. 6 shows an example of the variation of the C/I values for four different UEs over a series of 5000 TTIs. This plot was obtained by a simulation. As illustrated, for a given UE the range of C/I values may be as much as from around +12dB to -15dB. The C/I value varies due to shadowing, Rayleigh fading, and change in distribution of the mobile UEs, as well as cellular area specifications including the propagation parameters and speeds of UEs.

Fig. 7 is a graph illustrating a relationship between a data transmission rate and a channel-to-interference ratio (C/I) for four different modulation and coding combinations. The first three combinations

(are all quadrature amplitude modulation (QAM) schemes which differ from one another in the number (64 or 16) of constellation points used. The fourth combination uses quadrature phase shift keying (QPSK) as its
5 modulation scheme.

Each combination uses coding defined by a coding parameter which, in this example, is expressed as a redundancy rate R. For the first two combinations the redundancy rate R is 3/4, and for the third and fourth
10 combinations the redundancy rate is 1/2.

As can be seen from Figure 7, for C/I values lower than around -8dB the fourth modulation and coding combination (QPSK, R=1/2) is the only possible option. The characteristic of this combination is plotted with
15 crosses in the figure.

For C/I values in the range from around -8dB to around -2dB, the third combination (16QAM, R=1/2) provides the best transmission rate. The characteristic for this combination is illustrated by
20 triangular points in the figure.

For C/I values between around -2dB and +4dB the second combination (16 QAM, R=3/4) provides the best transmission rate. The characteristic of this combination is illustrated by square points in the
25 figure.

Finally, for C/I values greater than around +4dB, the first combination (64 QAM, R=3/4) provides the best transmission rate. The characteristic of this combination is illustrated by round points in the
30 figure.

The different combinations in Fig. 7 may be referred to as different MCS levels.

In the HSDPA system a technique such as adaptive modulation and coding (AMC) is used to adapt the MCS
35 level in accordance with the variations of the channel condition (e.g. C/I value). Each UE produces a measure

of the C/I of a downlink channel from the base station, and reports this measure (C/I value) to the base station. The base station then employs the reported C/I values for each UE, as well as information relating 5 to the system limitations and available MCS levels, to identify the most efficient MCS level for the particular UE. Thus, UEs that have better channels or are located in the vicinity of the base station can employ higher levels of MCS and therefore enjoy higher 10 transmission rates. The selection can be carried out, for example, by imposing thresholds (e.g. -8dB, -2dB, +4dB in the example illustrated with reference to Fig. 7) for moving to the next MCS level. Effectively, the result is a classification of the transmission rates 15 based on the channel quality of each UE.

Ideally, each UE reports a C/I value in every TTI and the base station is capable of setting a new MCS level for each available channel in every TTI. This means that each candidate solution η should preferably 20 include information specifying the MCS level to be applied to each available channel in the TTI under consideration. This MCS level is determined based on the latest reported C/I value for the receiver which, according to the candidate solution η concerned, is to 25 receive a packet in the TTI under consideration.

AMC enables coarse selection of the transmission rate in the HSDPA system. Another control technique which provides for a fine tuning ability of the data rate based on the channel conditions is referred to as 30 a hybrid automatic repeat request (H-ARQ) technique.

Fig. 8 is a schematic diagram for use in explaining how the H-ARQ technique works. In this example, the technique is a so-called stop-and-wait (SAW) version of the technique. The figure shows 35 packet transmissions in a single downlink channel HSPDSCH1 over a series of successive TTIs, TTI1 to

(
TTI9. In TTI2 a first packet is transmitted to UE1. Upon receiving a packet, each UE checks whether the transmission was error-free. If so, the UE sends an acknowledge message ACK back to the base station using an uplink control channel such as the dedicated physical control channel (DPCCH). If there was an error in the transmission of the received packet, the UE sends a non-acknowledge message NACK back to the base station using the uplink channel.

10 In the example shown in Fig. 8, the first packet transmitted to UE1 in TTI2 fails to be received error-free, and accordingly some time later, in TTI4, UE1 sends the NACK message to the base station. In the H-ARQ technique it is permitted for the next packet 15 destined for a particular UE to be transmitted without waiting for the acknowledge or non-acknowledge message of a packet previously transmitted to the same UE. Thus, none of the transmission timeslots can go idle in the case of error-free channels, which gives the 20 ability to schedule UEs freely. System capacity is saved while the overall performance of the system in terms of delivered data is improved.

For example, as shown in Fig. 8, before the NACK message for the first packet of UE1 is received by the 25 base station, the base station transmits a second packet to UE1 in TTI4. Thus, this second packet for UE1 is transmitted before the first packet for UE1 is retransmitted in TTI7 in response to the NACK message for the first transmission of the first packet.

30 In the H-ARQ technique, an erroneously-received packet (failed packet) is subject to a so-called chase combining process. In this process a failed packet is resent by the transmitter and subsequently the receiver 35 "soft" combines (for example using maximal ratio combining) all received copies of the same packet. The final carrier-to-interference ratio (C/I) is determined

as the sum of the respective SIRs of the two packets being combined. Thus, the chase combining process improves the C/I of the transmitted packets.

All retransmissions have a higher priority than 5 first transmissions. This means that all retransmission packets are given the opportunity to be transmitted before the first transmissions of new packets.

The HSDPA system specifies the number of 10 retransmissions permitted. If a packet cannot be delivered error-free within the permitted number of retransmissions, it will be dropped. The number of permitted retransmissions is a parameter that has a high impact on the overall system performance. 15 Increasing the number of retransmissions results in a better C/I which in turn improves the frame error-rate (FER). However, since retransmissions have higher priority than first transmissions, increasing the permitted number of retransmissions will also increase 20 the delivery delay and have a negative impact on the throughput. Accordingly, in the HSDPA system a trade-off exists between the number of permitted retransmissions, the throughput and the delay of the system. Incidentally, the MCS level for first 25 transmissions and retransmissions may be different.

The genetic algorithm unit 128 is represented schematically in Fig. 9. It makes use of a plurality (n) of individuals 30_1 to 30_n . n is, for example, 100. As described later in more detail, each individual has 30 its own set of parameters (UE vector) representing a candidate scheduling solution for the scheduling instant under consideration.

At the end of the processing for the scheduling instant under consideration, only one of the available 35 individuals is selected as the best solution. However, all the individuals are subject to processing by a

genetic algorithm over a series of iterations (generations). This is illustrated schematically in Fig. 9 which shows that in each iteration all of the individuals are ranked for fitness by a fitness function and then subject to selection, crossover and mutation as part of the genetic algorithm, with the individual 30_x being selected as the best solution at the end of the series of iterations.

Before describing the detailed operation of the genetic algorithm unit 128 in Fig. 10, some background information regarding genetic algorithms will be provided.

When considering how to apply a genetic algorithm to a particular problem, four steps are required. Firstly, a way of representing the solution must be chosen. Secondly, a random variation operator must be devised. Thirdly a rule for solution survival must be determined. Fourthly, the population must be initialised.

Dealing with these steps in more detail, to represent any possible solution within the confines of the genetic algorithm, a structure must be defined for the data that can encode every possible solution that it might be desirable to evaluate (step 1). There is no single best choice for the representation. Also, the cost function, i.e. the means to evaluate a fitness value for any candidate solution, must also be determined. The level of complexity of the cost function is related to the amount of knowledge involved. In terms of complexity of the evaluation of the cost function a genetic algorithm is more attractive than classic or neural-network-based techniques. In a genetic algorithm the fitness value does not need to be a precise quantity. It can be any rough calculation that ranks the individuals in the right order in terms of their performance. Having a

simple but efficient fitness function leads to a considerable reduction in the complexity of the algorithm.

Many options exist for devising the random variation operator (or operators) that can be used to generate child solutions from parent solutions (step 2). In nature, there are two general forms of reproduction: sexual and asexual. In sexual reproduction, two parents within a species exchange genetic material that is recombined to form a child. Asexual reproduction is essentially cloning, but mutations of various forms can creep into the genetic information passed from parent to child. These operators are worth modelling in the genetic algorithm. However, it is not necessary to be limited to random variation operators found in nature. As will be explained in more detail later, the solutions may be influenced by a deterministic algorithm as they evolve from one generation to the next. Also, genetic material can be recombined from three or more parents with some kind of democratic vote between the parents involved in reproduction.

There is virtually no limit to the types of variation operators that can be devised. The ultimate success of a genetic algorithm depends strongly on how well the variation operators, the representation and the fitness function are matched. Different operators will vary in usefulness according to the situation. It is provable mathematically that, just as with the solution representation, there is no single best variation operator for all problems.

The rule for solution survival (step 3) is also referred to as the selection operator, as it selects which solutions will survive to become the parents of the next generation. As with the other items, many forms of selection can be considered. One simple rule

(
is the survival of the fittest. Only a handful of the very best solutions in the population are retained, while all the other solutions are killed off. An alternative is to use a sort of tournament approach, 5 where randomly-paired solutions compete for survival. Just as in a tournament, where weaker players sometimes win through the first few rounds because they get a series of lucky draws in the tournament, weaker solutions in a population in a genetic algorithm 10 sometimes survive a few generations with this approach. This can be an advantage in complex problems, where it may be easier to find new improved solutions by making variations of weaker ones than to do so by relying only on the very best. The possibilities are plentiful, but 15 any rule that generally favours better solutions over weaker solutions for survival is reasonable.

Initialising the population (step 4) depends on what, if any, knowledge exists at the outset about how to solve the problem. If nothing is known about how to 20 solve the problem, then solutions can be chosen completely randomly from the space of all the possible solutions. However, if there is any extra knowledge available, it can be used in creating the initial population. That said, while it is possible to 25 incorporate any problem-specific knowledge available and thereby take advantage of it when using genetic algorithms, it is not inherently necessary for the genetic algorithm to succeed. This is why genetic algorithms can tackle an enormously broad range of 30 problems.

Next, an embodiment of the present invention suitable for use in an HSDPA system will be described in detail. This embodiment seeks to optimise the performance of the HSDPA system in terms of measures 35 such as quality of service (QoS), fairness of packet delivery, total delay, total throughput and average

throughput. The optimum channel assignment and packet scheduling policy can be defined as

$$\max_{\substack{(\overline{Oct})_k, (\overline{MCS})_k, (\overline{Po})_k, (\overline{UE})_k \\ k=1 \dots \alpha}} f(Th, Fa, 1/De, QoS, Avg_Th) \Big|_{T, (\overline{C/I})_k}, \quad (1)$$

5 where $f(\bullet)$ is a function (fitness function) that combines all the desired performance-related measures to a unique value, as described later in more detail.

Here α is the number of TTIs.

10 Vector $(\overline{Oct})_k$ is the decision for the number of octets to be transmitted in the k th TTI on each of C channels so that

$$\overline{Oct} = \{Oct_1, Oct_2, \dots, Oct_C\}, \quad (2)$$

15 where Oct_i is the number of octets assigned to channel i .

20 Similarly, vector $(\overline{MCS})_k$ defines the decision on the MCS levels for the k th TTI for the C channels so that

$$\overline{MCS} = \{MCS_1, MCS_2, \dots, MCS_C\}, \quad (3)$$

where MCS_i is the MCS level assigned to channel i .

25 vector $(\overline{Po})_k$ is defined as the power allocated to each channel so that

$$\overline{Po} = \{Po_1, Po_2, \dots, Po_C\}, \quad (4)$$

30 where Po_i is the power allocated to channel i .

(

Vector $(\overline{UE})_k$ is a UE vector which defines the decision on the UE which will be assigned to which channel so that the vector \overline{UE} is

5

$$\overline{UE} = \{UE_1, UE_2, \dots, UE_C\}. \quad (5)$$

where UE_i is the candidate UE to which data is to be delivered on channel i .

10

Vector $(\overline{C/I})_k$ is employed to represent the reported C/I values on each channel for the k th TTI

$$\overline{C/I} = \{C/I_1, C/I_2, \dots, C/I_C\}. \quad (6)$$

15

where C/I_i is the reported C/I value on channel i .

The total throughput of the HSDPA system Th in (1) may be defined as

$$Th = \frac{\sum_{n=1}^F (Oct_{Received})_{UE_n}}{\sum_{n=1}^F (Oct_{Arrived_Node_B})_{UE_n}}, \quad (7)$$

20

where F is the number of UEs with active sessions, $(Oct_{Received})_{UE_n}$ is the number of successfully delivered octets to the n th UE and $(Oct_{Arrived_Node_B})_{UE_n}$ is the number of octets originally delivered to the source queue of the n th UE at the transmitter (Node B).

25

The individual throughput of each UE is defined as

$$Th_n = \frac{(Oct_{Received})_{UE_n}}{(Oct_{Arrived_Node_B})_{UE_n}}, \quad n = 1 \dots F, \quad (8)$$

and the vector of the individual throughputs as

$$\overline{Th} = \{Th_n\}, \quad n = 1 \dots F, \quad (9)$$

The mean of the individual throughputs or Avg_Th may be
5 defined as

$$Avg_Th = mean(\overline{Th}), \quad (10)$$

and Fa or fairness as the variance of \overline{Th}

$$Fa = var(\overline{Th}), \quad (11)$$

10 where $var(\bullet)$ represents the variance operation.

To measure the delivery delay, each octet of
packet data may be stamped on its arrival at the
transmitter (Node B). The time at which it is correctly
15 delivered to the UE is recorded. The delivery delay is
defined as the difference between the arrival time at
the Node B and the successful delivery. This delay may
be defined as

$$delay_i = (delivery_time_i)_{UE} - (arrival_time_i)_{Node_B}, \quad (12)$$

20 In (1) the total delay De for the HSDPA system may be
defined as

$$De = \sum_{i=1}^{\theta} delay_i, \quad (13)$$

25 where θ is the total number of the received octets for
all UEs.

A *QoS* parameter in (1) may be defined as the ratio
of the successfully delivered data in a defined
threshold of time γ to the total number of transmitted
30 octets. This defined threshold of time γ is preferably
different for different types of data. For example,

(
 voice or video data will have a lower threshold (higher priority) than an e-mail message or WWW download. The delays calculated in (12) are preferably also expressed relative to the threshold γ applicable to the data type.

5 A preferred embodiment of the present invention will now be described in more detail with reference to Figs. 9 to 17 and equations (14) to (38).

10 In this embodiment it is assumed that the allocated powers for the channels are equal, i.e.

$$Po_1 = Po_2 = \dots = Po_C, \quad (14)$$

Looking at the optimum scheduler and channel assignment in (1) it can be seen that the optimum assignment requires three vectors $(\overline{Oct})_k$, $(\overline{MCS})_k$ and $(\overline{UE})_k$ with length α . It is preferable to try to construct these vectors as time progresses for each TTI. Thus, in this embodiment each candidate solution is represented as

20 $\eta = \{\overline{UE}, \overline{C/I}, \overline{Oct}, \overline{MCS}\}, \quad (15)$

The MCS levels, \overline{MCS} , are decided based on the C/I reported values $\overline{C/I}$. Finally depending on \overline{MCS} levels and the transmission or retransmission states a decision is made on the number of available octets per channel $(\overline{Oct})_k$.

25 Fig. 10 is a flowchart for use in explaining operation of the first embodiment. The processing shown in Fig. 10 is performed for each scheduling instant under consideration. Each scheduling instant corresponds to one transmission time interval (TTI), for example. Of course, the scheduling required for a particular scheduling instant (TTI) must be determined

in advance of the time at which that scheduling instant actually occurs, because the selected scheduling solution η_{best} is needed at the start of the TTI concerned.

5 As indicated earlier, each candidate solution generated by the genetic algorithm unit 138 is represented by an individual.

In step S2 at the beginning of the processing for the scheduling instant under consideration an 10 initialisation process is carried out to create the first generation of individuals or population. Each population of individuals is defined as

$$\Pi_\rho = \{\beta_j\}_\rho, \quad j=1 \dots n \quad (16)$$

15 where Π_ρ is the population of individuals of generation ρ , n is the population size, β_j is the chromosome of the individual and j represents the index of each individual.

The structure of one individual is illustrated 20 schematically in Fig. 11. The individual 30 has a chromosome 32 (β). In Figure 11, t represents the chromosome length so that

$$t = C \quad (17)$$

25 where C is the number of available channels.

In a CDMA system, each channel has an unique channelisation code.

The chromosome 32 is made up of a plurality of 30 individual genes α_1 to α_c representing the UEs to be sent packets in the scheduling instant under consideration. Thus, α_1 specifies the UE to be assigned to channel 1, α_2 specifies the UE to be assigned to channel 2, and so on.

(The individual 30 also has a set of further elements 34 to 42 holding parameters used by the genetic algorithm. The parent 1 and parent 2 parameters in elements 34 and 36 are index values used 5 to identify first and second parents of the individual 30. The fitness value in element 38 is a parameter used for ranking the individuals in the population of a generation. The crossover site 1 and crossover site 2 parameters in elements 40 and 42 are used to identify 10 first and second positions along the chromosome 32 for crossover operations of the genetic algorithm. These first and second positions are referred to as crossover site 1 and crossover site 2 respectively.

15 The first generator 122 in the candidate solution generating unit 12 is activated to produce a max C/I solution $\eta_{C/I}$. This solution is assigned to the chromosome β_1 of the first individual 30_1 .

$$\eta_{C/I} \rightarrow \beta_1 \quad (18)$$

20 The second generator 124 in the candidate solution generating unit 12 is activated to produce a round robin solution η_{RR} . This solution is assigned to the second individual 30_2 ,

$$\eta_{RR} \rightarrow \beta_2 \quad (19)$$

25 The third generator 126 is the candidate solution generating unit 12 is activated to produce a plurality of random solutions η_{RM} . These solutions are assigned respectively to the remaining individuals 30_3 to 30_n .

$$30 \quad \eta_{RM} \rightarrow \beta_j, \quad j = 3 \dots n \quad (20)$$

Next, in step S3 a set of individual fitness measures F is produced for each candidate solution η_i to

η_n) in the current population. The fitness measure producing unit 142 in the solution selecting unit 14 may be used for this purpose.

5 In this embodiment the fitness of each solution is evaluated by estimating the contribution of the solution to throughput, delay and fairness.

As the first fitness measure the unit 142 estimates the number of the delivered octets if the current solution η is chosen so that

10

$$Eff_Oct(\eta) = \sum_{i=1}^C \overline{Oct}(i).(1 - FER(UE(i))), \quad (21)$$

where $UE(i)$ is the UE assigned to channel i and $FER(UE(i))$ is the estimated FER for channel i . When the frame error rate in a channel increases the possibility of receiving error-free octets decreases.

15 As the second fitness measure for each solution η , in the current TTI and for the candidate UEs , the unit 142 measures the ratio of octets waiting at the transmitter octets for transmission to the total number 20 of octets arrived in the Node B source queue of that UE so far, so that

$$\begin{aligned} Ratio_Waiting_Oct(\eta) = & \\ \sum_{i=1}^{\omega} Waiting_Oct(Selected_UE(i)) / Arrived_Oct(Selected_UE(i)) & \end{aligned} \quad (22)$$

25 where $Selected_UE$ is the set of the selected UEs in solution η and ω is the number of candidate UEs in the current solution so that

$$1 \leq \omega \leq C \quad (23)$$

30 As the third fitness measure the unit 142 measures the delay experienced by the earliest undelivered octet

(which has arrived at the transmitter but has not yet been delivered. This measure is defined as

$$\text{Delay_Profile}(\eta) = \frac{1}{\sum_{i=1}^C (N.TTI - \text{Arrival_Time_Earliest}(UE(i)))} \quad (24)$$

5 where $N.TTI$ represents the time at the current TTI, $\text{Arrival_Time_Earliest}()$ is a vector which includes the arrival time of the earliest undelivered octet in Node B for each UE , and $UE(i)$ is the UE assigned to channel i for the current solution η .

10 Finally the fourth fitness measure is defined based on the fairness of the channel assignment and scheduling process for the HSDPA system. The fourth fitness measure is effectively an estimate of the impact of the candidate solution on the fairness of 15 packet scheduling process. For the current TTI, the number of successfully received octets with ACK messages is determined. It is defined as

$$\text{Received} = \{R_i\}_{i=1 \dots F} \quad (25)$$

20 where F is the number of UE s with active transmission sessions, i is the UE number and R_i is the number of successfully received octets for i th UE . The estimated received octets for the current solution η is defined as

25

$$\text{Estimated_Received}(\eta) = \{\bar{R}_i\}_{i=1 \dots F} \quad (26)$$

where F is the number of UE s with active transmission sessions, i is the UE number and \bar{R}_i is the number of successfully received octets for i th UE so that for 30 each solution η

$$\hat{R}_{Selected_UE(c)} = R_{Selected_UE(c)} + \overline{Oct}(c), \quad c = 1 \cdots C, \quad (27)$$

where c indicates the channel number. Based on the estimated number of received octets, for the current solution η and the current TTI, the expected throughput is estimated for all F active UEs so that

$$\hat{Th}(\eta) = \left\{ \frac{\hat{R}_i}{Arrived_Oct(i)} \right\}_{i=1 \cdots F}, \quad (28)$$

The expected contribution of the current solution to the fairness of the packet scheduling process is defined as

$$Fairness(\eta) = \frac{1}{\text{var}(\hat{Th}(\eta))}, \quad (29)$$

Based on the individual fitness measures produced for each candidate solution a unified measure of fitness is calculated for each candidate solution, and the respective unified measures for the different candidate solutions are compared to select the parent individuals of the current generation.

Before describing the process performed in the present embodiment to select parent individuals, an explanation of a preferred unified fitness measure will be described with reference to Figs. 12 to 16. This preferred unified measure of fitness in the present embodiment is based on a product of the individual fitness measures for the candidate solution concerned. To facilitate the explanation, in these figures only three of the four fitness measures are represented, namely the first, third and fourth measures described above. This enables each set of fitness measures to be represented as a plane in three-dimensional space

(having dimensions of Eff_Oct, Delay_Profile and Fairness.

In Fig. 12, a first candidate solution η_1 has a first set of values E_1 , D_1 and F_1 for the three fitness measures. A second candidate solution η_2 has a second set of fitness values E_2 , D_2 and F_2 . In the case illustrated in Fig. 10, $E_1 > E_2$, $D_1 > D_2$ and $F_1 > F_2$ and it is absolutely clear that the solution η_1 having the higher plane is the better solution.

A more complicated scenario is illustrated in Fig. 13. In this case, if the unified fitness measure is based on a product of the individual fitness measures, there is a risk that one solution which has a much higher value of a particular fitness measure than all the other solutions will dominate, even if in terms of the remaining fitness values it is a worse solution. As shown in Fig. 13, the solution η_1 has a much higher fairness value F_1 than that of the solution η_2 . The remaining fitness values for solution η_1 are worse than for the solution η_2 . However, the product $F_1 \cdot E_1 \cdot D_1$ of the fitness measures for solution η_1 is greater than the product $F_2 \cdot E_2 \cdot D_2$ of the fitness measures for solution η_2 , so the solution η_1 dominates. Although a situation such as that shown in Fig. 13 can always arise, there is less chance of it arising as the number of candidate solutions increases. Accordingly, in embodiments of the present invention it is desirable to have as many iterations per scheduling instant as can be performed within the available processing time.

Fig. 14 shows another problem (mismatch) which needs to be addressed when designing the unified fitness measure. In this figure, even for any particular solution one of the fitness measures has a much higher value than each of the other fitness measures. In this case, any product-based unified measure will place undue emphasis on the single fitness

measure having the highest value, with the risk that performance in terms of the other fitness measures is neglected.

In order to try to address this mismatch problem,
 5 it is preferable to apply a mapping function to one or more of the individual fitness measures so that after mapping the fitness measures are of the same order as illustrated in Fig. 15. The mapping function is, for example, a normalising function. For example, each
 10 fitness measure may be scaled relative to a peak value of the measure identified from previous outcomes.

It is also preferable that the unified fitness measure be controllable or tunable to create different scenarios. This can be achieved by weighting the
 15 different individual fitness measures of a candidate solution using controllable weighting coefficients. An extreme example of such a weighting process is illustrated in Fig. 16. In this example, the throughput measure Eff_Oct has a weighting coefficient of zero, with the result that the throughput fitness
 20 measures of all candidate solutions have the same value, for example one. The remaining fitness measures have non-zero weighting coefficients so that these other measures exclusively influence the unified fitness measure. Accordingly, the solutions with maximum fairness and minimum delay will tend to be
 25 selected in this scenario.

Taking the matters set out above into account, a preferred function for determining the unified measure of fitness, for the candidate solution η , i.e. fitness function $f(\eta)$, is

$$f(\eta) = \frac{W_{1E} + W_{2E} \cdot u(Eff_Oct)}{W_{1E}} \cdot \frac{W_{1D} + W_{2D} \cdot v(Delay_Profile)}{W_{1D}} \cdot \frac{W_{1F} + W_{2F} \cdot x(Fairness)}{W_{1F}} \cdot \frac{W_{1R} + W_{2R} \cdot y(Ratio_Waiting_Oct)}{W_{1R}} \quad (30)$$

(
 where $u(\cdot)$, $v(\cdot)$, $x(\cdot)$ and $y(\cdot)$ are the mapping functions, W_{1E} , W_{2E} , W_{1F} , W_{2F} , W_{1D} , W_{2D} , W_{1R} and W_{2R} are the weighting coefficients.

5 The present embodiment therefore tries to maximise the unified fitness measure considering all the QoS parameters in (1). By employing a tuning mechanism, it reaches a high level of flexibility and control over different parameters such as fairness, total
 10 throughput, average throughput and total delay.

Conventional scheduling techniques such as RR or Max C/I in a way try to perform this maximisation process per TTI. However RR scheduling, for example, puts most effort on Fa in (1) by ignoring the other
 15 parameters such as Th to some extent while max C/I scheduling mostly concentrates on Th and De .

Initially a decision is made, based on the selected scheduling policy, about the values to be assigned to the weighting coefficients. For example in
 20 (30), if just the fairness matters then W_{1F} and W_{2F} are each set to one, W_{2E} , W_{2D} , W_{2R} are set to small values and W_{1E} , W_{1D} , W_{1R} are set to very large values. The algorithm then tends to behave similar to RR.

Referring back now to Fig. 9, after producing a
 25 unified fitness measure $f(\eta)$ for each candidate solution (i.e. each individual) in step S3, a selection process for selecting parent individuals based on the unified fitness measures is carried out in step S4. In step S4, parents to be used to create the next
 30 generation ρ of individuals are selected from the individuals of the current generation $\rho-1$, in accordance with a so-called random roulette wheel selection process so that

$$\left(\left(\beta_j \right)_{\text{parent}} \right)_\rho = S(\Pi_{\rho-1}), \quad j=1 \cdots n \quad (31)$$

where $S(\cdot)$ is the selection process and $((\beta_i)_{parent})_p$ represents j th parent of the p th generation.

The selection process gives better survival probabilities to fitter individuals than to their weaker counterparts. Assume that $(f_{sum})_{p-1}$ is the sum of the fitness measures of the individuals of the current generation. Also assume that $(f_j)_{p-1}$ is the fitness measure of the individual j of the current generation. The probability of any individual to be selected from the population may be defined as

$$P_s(j) = (f_j)_{p-1} / (f_{sum})_{p-1} \quad (32)$$

where $P_s(j)$ is the probability of the selection of individual j of the current generation. The individual with the index j is selected if

$$\sum_{i=1}^j (f_i)_{p-1} < (f_{sum})_{p-1} z, \quad j = 1 \dots n, \quad (33)$$

where z is a real random number between 0 and 1. The highest value of j only that satisfies eqn. 33 is selected. This selection process, $S(\cdot)$, is a linear search through a roulette wheel. This roulette wheel is of circumference f_{sum} and is segmented into n slots of different lengths corresponding to the fitness measures f_j of the individuals. The real random value z in (33) indicates the place where the wheel has stopped after a random spin. To perform a linear search, it is necessary to have a uniform distribution for the probability density function PDF of z .

In this embodiment, it is assumed that parents of the current generation will be selected in pairs, and each pair of selected parents will produce two child individuals which will become part of the population for the next generation. The individuals of the current generation, whether selected as parents or not,

(will be killed off when the current generation ends and the next generation begins. To maintain the population at a constant level n from one generation to the next, $n/2$ pairs of parents are selected in each generation to be parents. Incidentally, in the present embodiment, the same individual may be selected more than once as a parent, i.e. bigamy is permitted. It is also possible, however, to design the selection process in such a way that an individual can only be selected once as a parent, if desired.

Next, in step S5 the pairs of selected parents are subjected to crossover and mutation processes to form first and second child individuals for each pair of parents. For each pair of selected parents the crossover sites are defined. The crossover sites are chosen based on a crossover probability. These sites define the amount of the genetic material from the two parents that should be recombined and transferred into their children. In this embodiment, a two-point crossover is used to select two sites (positions) X_1 and X_2 to cut the chromosomes.

In the crossover process the parent genes between the crossover sites X_1 and X_2 are exchanged, i.e. the genes of parent 1 are transferred to (inherited by) child 2 and the genes of parent 2 are transferred to (inherited by) child 1.

Outside the two crossover sites no exchange of parent genes is carried out, i.e. the genes of parent 1 are transferred to (inherited by) child 1, and the genes of parent 2 are transferred to (inherited by) child 2.

In this embodiment, the child genes outside the crossover sites are, however, subject to a mutation process which brings about a random change in a child gene with a probability called the probability of mutation. The maximum amount of the random change may

be set by a mutation range parameter ε . In this embodiment, ε is a variable parameter set equal to the number of UEs which at any given time have non-empty source queues at the transmitter, i.e. have data waiting at the transmitter for transmission.

Thus, as shown in Fig. 17, in step S5 a joint crossover and mutation process $CM()$ is carried out to create chromosomes β_j and β_{j+1} of first and second child individuals from the chromosomes β_E and β_F of a pair of first and second parent individuals:

$$\left(\left(\beta_j, \beta_{j+1} \right)_{\text{Child}} \right)_\rho = CM \left(\left(\beta_E, \beta_F \right)_{\text{Parent}} \right)_{\rho-1}, P_x, P_m, X_1, X_2, \varepsilon. \quad (34)$$

where P_x is the crossover probability, P_m is the probability of mutation, X_1 and X_2 are the crossover sites, E and F are the indexes of the selected parents of the individuals belonging to the current generation and ε is the mutation range parameter. Each crossover point X_1 or X_2 is calculated as

20

$$\begin{cases} X = \text{int}(z\ell) & \text{if } \text{coin}(P_x) = T \\ X = \ell & \text{if } \text{coin}(P_x) = F \end{cases} \quad (35)$$

where $\text{int}()$ calculates the integer part, z is a real random number between 0 and 1; $\text{coin}(P_x)$ is the process of the tossing of a biased unfair coin with the crossover probability P_x and ℓ is the chromosome length, T is true (e.g. heads) and F is false (e.g. tails).

Equation (35) is applied twice to produce two values for X . The smaller value becomes X_1 and the larger value becomes X_2 .

$$\begin{cases} \mathbf{C}_j \mapsto \mathbf{a}_i = \mathbf{P}_E \mapsto \mathbf{a}_i, & 1 \leq i \leq X_1 \\ \mathbf{C}_j \mapsto \mathbf{a}_i = \mathbf{P}_F \mapsto \mathbf{a}_i, & X_1 < i < X_2, \\ \mathbf{C}_j \mapsto \mathbf{a}_i = \mathbf{P}_E \mapsto \mathbf{a}_i, & X_2 \leq i \leq t \end{cases} \quad (36)$$

and

$$\begin{cases} \mathbf{C}_{j+1} \mapsto \mathbf{a}_i = \mathbf{P}_F \mapsto \mathbf{a}_i, & 1 \leq i \leq X_1 \\ \mathbf{C}_{j+1} \mapsto \mathbf{a}_i = \mathbf{P}_E \mapsto \mathbf{a}_i, & X_1 < i < X_2, \\ \mathbf{C}_{j+1} \mapsto \mathbf{a}_i = \mathbf{P}_F \mapsto \mathbf{a}_i, & X_2 \leq i \leq t \end{cases} \quad (37)$$

5

where E and f are the two selected parents, P represents the current generation, C represents the next generation, X_1 and X_2 are the crossover sites and two new born children have the indexes j and $j+1$.

10 In the mutation process each gene outside the crossover sites is mutated in a random process so that

$$\begin{cases} \alpha_{Child} = UE_full_Queue(1 + \text{int}(z\epsilon)) & \text{if } coin(P_m) = T \\ \alpha_{Child} = \alpha_{Parent} & \text{if } coin(P_m) = F \end{cases}, \quad (38)$$

15 where $UE_full_Queue()$ represents the set of UEs with non-empty source queues at Node B which expect data delivery, ϵ is the variable mutation range parameter which is equal to the number of UEs in the set, $coin(P_m)$ is the process of the dropping of a biased unfair coin with mutation probability P_m , z is a real uniform random number between 0.0 and 1.0 and $\text{int}()$ returns the integer part of a real number. In this way, the mutated gene is set to be a random one of the UEs having data waiting at the transmitter.

20 In step S6 it is determined whether or not enough generations of individuals have yet been processed. It is possible for the number of generations required to be a fixed threshold value, for example 50 to 100. The

number of generations required is dependent on the population size. As the population size increases the number of generations required to reach an acceptable solution decreases. Alternatively, in step S6 it could 5 be decided that sufficient generations have been processed if the best fitness measure amongst the child individuals created in step S5 exceeds a desired threshold value, indicating that a suitable candidate solution has already been found.

10 If the outcome in step S6 is that sufficient generations have not yet been processed, then in step S7 the current population is killed off in this embodiment. The new population for the next generation is made up entirely of the child individuals 15 created in step S5. Processing then returns to step S3 for the next generation.

If in step S6 it is determined that sufficient 20 generations have now been processed, in step S8 the current generation becomes the last generation and the child individual of that last generation which has the best fitness measure is selected as the best candidate 25 solution η_{best} by the solution selecting unit 14 (Fig. 5). Processing is then terminated in step S9.

Fig. 18 is a flowchart for use in explaining the 25 processing carried out by the second generator 124 in Fig. 4.

In a first step S20 processing starts for the scheduling instant (TTI) under consideration.

The generator 124 maintains a list of UEs to which 30 data is waiting at the transmitter to be transmitted. A UE has waiting data if its source queue is non-empty. An example of the waiting list is shown in Fig. 19(a). In step S21 the UE at the bottom of the list is moved to the top. In the present example, UE34 down at the 35 bottom of the list in Fig. 19(a) is moved to the top of

the list. All other UEs are moved down one place in the list.

In step S22 the list is updated by adding any new UEs which now have waiting data and deleting any existing UEs for which there is no longer any waiting data. Thus, after steps S21 and S22 the list may appear as shown in Fig. 19(b). UE49 has been deleted as it no longer has any waiting data, and a new UE, UE30, is added to the bottom of the list.

In the simplest RR solution the UE with the highest priority (UE34 in Fig. 19(b)) is allocated all the available channels (channels 1 to 10). In step S23 a measure of throughput is calculated for each available MCS level (and at the particular power level chosen). Also, a C/I value is obtained from a C/I report received from the specified UE (i.e. UE34 in this example). Then in step S24 the generator 124 decides the most efficient MCS level for each channel. After this, the amount of data (number of octets) to be transmitted on each channel is decided.

In step S25 all the available channels are assigned to the selected UE (the UE at the top of the list).

In step S26 the processing for the scheduling instant under consideration is terminated and the solution η_{RR} is output.

Fig. 20 shows a variation of the basic RR scheduling process which is referred to as a packet-weighted RR process. Steps S20 to S22 are the same as in Fig. 18.

In step S30, for each UE in the list, a measure of throughput is calculated for each available MCS level. Also, a C/I report is obtained from each UE in the list.

In step S31 the UE at the top of the list is selected as the current UE to be processed. In this

case, however, the UE at the top of the list is not the only UE which will be allocated channels. Thus, an assessment is made in step S31 of the source queue fill level for the current UE in relation to the other UEs in the list. Then, in step S31, the number of channels to be allocated to the current UE, the MCS level and the number of octets to be transmitted on each channel allocated to the current UE are decided.

In step S32 the determined number of channels is assigned to the selected user. For example in Fig. 21(a) the UE11 is assigned channels 1 and 2. The minimum number of channels assignable is 1. In step S33 the number of channels remaining for allocation is determined. In this example, after allocating channels 1 and 2 to UE11, eight channels (channels 3 to 10) remain.

In step S34 it is checked whether there is any remaining available channel. If there is, then in step S36 the next UE in the list is selected (UE5). The process of steps S31 to S34 is repeated for this UE. Thus, as shown in Fig. 21(a) UE5 is allocated channel 3. The processing continues in this way until all the available channels have been allocated. Eventually, after allocating the last channel (channel 10) to UE8, there is no channel remaining. The remaining UEs do not get allocated channels in this scheduling instant. In this example, UEs 23 and 34 are not allocated channels.

Fig. 22 shows an example of the process performed by the first generator 122 in Fig. 4 to generate the max C/I candidate solution $\eta_{c/i}$.

In step S20 processing begins for the scheduling instant under consideration. In step S40 the list of UEs with waiting data is updated in the same way as in the previously-described step S22 of the RR process.

After this, in step S42, the list is reordered in order of decreasing reported C/I value.

The following steps S44 to S47 correspond respectively to the previously-described steps S23 to 5 S26 in the RR process of Fig. 18. Thus, in this process all the channels are allocated to just the single UE at the top of the list, i.e. the UE with the best reported C/I value.

It is also possible for the first generator to use 10 a packet-weighted variant of the process, as shown in Fig. 23.

This process is similar to the basic max C/I process of Fig. 22. However, instead of selecting a 15 single UE at the top of the list and allocating all the available channels to that one UE, the packet-weighted process allocates channels to two or more UEs occupying the highest positions in the list. This is done by first applying the step S50 and then applying steps S51 to S56 iteratively. Steps S50 to 20 S56 are identical respectively to steps S30 to S36 in the packet-weighted RR process of Fig. 20.

Incidentally, as is known in the wireless communication art, it is possible for a UE to have two or more possible channels available for use in the 25 downlink direction. These channels may have different carrier frequencies, for example. In this case, the UE produces C/I reports for each available channel, and the base station and/or the mobile station selects one of the available channels (carriers). This may be 30 taken into account by a scheduler embodying the present invention, so that the selection of the available channel is part of each candidate solution.

Next, simulation results for embodiments of the 35 invention are provided. These simulation results relate to simulated worldwide web (WWW) browsing sessions. The simulations are based on parameters

(described in TR25.848, ver1.0.0, RP-010191, TSG-RAN#11, March 2001. A WWW browsing session comprises a sequence of packet calls. A UE initiates a packet call when requesting an information entity. In the present simulation, a TTI is assumed to be of duration 2ms (as currently specified in the 3GPP standard). The concept of "efficient throughput" for a service is defined as the portion of the throughput which satisfies the QoS requirements within the specified delay threshold γ .
5
10 In the present simulation it is assumed that the acceptable delay threshold γ (delivery time delay) is 1.5 seconds. Chase combining is carried out and a maximum of six retransmissions is permitted. A packet is dropped if it cannot be delivered within six
15 retransmissions.

It is also assumed that the channel estimation performed by each UE is perfect, and that the feedback signalling is error-free. A minimum reporting delay is considered to be three TTIs. This represents the
20 minimum time between the transmission of a packet and the receipt of an acknowledgement message ACK from the UE. The simulation was carried out for a period of ten seconds (5000 TTIs and for 50 UEs).

Fig. 24 shows the simulated packet data arriving
25 in the source queues at the transmitter base station for the 50 UEs.

For each UE a simulated C/I scenario is generated, for example using a Gaussian autoregressive model. In order to simulate extreme conditions the UEs are assumed to have significant C/I differences. As shown
30 in Fig. 25, the C/I value is the same for any given UE over all TTIs. The fixed C/I value increases progressively from -2dB for UE50 to +5dB for UE1. Accordingly, the best C/I value belongs to UE1 and the
35 worst C/I value belongs to UE50.

Fig. 26 is a graph for comparing overall throughput versus delay in a scheduler embodying the present invention (solid line) and in a max C/I scheduler (dotted line) and in a round robin scheduler (dot-dash line). The lines indicate a cumulative density function (CDF) of the throughputs of all UEs. The results in Fig. 26 demonstrate that a scheduler embodying the present invention manages to outperform the max C/I scheduler and the round robin scheduler in terms of the overall delivered throughput.

Table 1 below presents the total delay (in seconds) and the average delay (also in seconds) of successfully-delivered packets in the round robin scheduler, the max C/I scheduler and a scheduler embodying the present invention. The total delay is the sum of the respective delays experienced by all successfully-delivered packets. The average delay is the average of the respective delays experienced by the successfully-delivered packets.

20

Scheduling Type	RR	Max. C/I	Invention
Total Delay (seconds)	79175.0	57184.0	41080.0
Average Delay (seconds)	2.0846	1.539	1.2608

Table 1

As is evident from Table 1 a scheduler embodying the present invention outperforms both the round robin scheduler and the max C/I scheduler in terms of delay profile. Significantly, whereas the average delay for both the round robin scheduler and max C/I scheduler exceeds the target delay value of 1.5 seconds, the average delay achieved by a scheduler embodying the

present invention is significantly lower than this target value.

The superior performance of a scheduler embodying the present invention in terms of QoS (efficient total throughput within the 1.5 second delay window) is illustrated in Fig. 27 which shows the first 5 seconds of delay in Fig. 26. In this case, however, two different schedulers embodying the present invention are illustrated. The first of these, labelled embodiment 1 and illustrated by a dashed line in Fig. 27, is the same as in Fig. 26. As is clear from Fig. 27, this embodiment manages to deliver almost 0.08 more packets within the acceptable delay threshold of 1.5 seconds than the max C/I scheduler.

As explained above, a scheduler embodying the present invention preferably has a tunable or controllable fitness function, for example the fitness function of equation (30). The second characteristic for a scheduler embodying the present invention in Fig. 27, labelled embodiment 2 and indicated by a dot-dash line, illustrates how this tuning can be used to influence the performance. In this second embodiment the tuning is changed compared to the first embodiment so as to increase the fairness of packet delivery and decrease the total throughput whilst maintaining the importance of delay unchanged. As shown in Fig. 27, this leads to a scheduler having better overall fairness than the max C/I scheduler and better total throughput than the round robin scheduler.

Accordingly, a scheduler embodying the present invention can outperform the conventional scheduling techniques such as max C/I and RR considerably. A scheduler embodying the present invention can also achieve a high level of tuning ability and flexibility to control various system parameters.

Although an example of the present invention has been described above in relation to a wideband CDMA network having an asynchronous packet mode, it will be appreciated that it can also be applied to any other networks in which scheduling is required. These networks could be, or could be adapted from, other CDMA networks such as an IS95 network. These networks could also be, or be adapted from other mobile communication networks not using CDMA, for example networks using one or more of the following multiple-access techniques: time-division multiple access (TDMA), wavelength-division multiple access (WDMA), frequency-division multiple access (FDMA) and space-division multiple access (SDMA).

Although embodiments of the present invention have been described as having distinct "units", those skilled in the art will appreciate that a microprocessor or digital signal processor (DSP) may be used in practice to implement some or all of the functions of the base station and/or mobile station in embodiments of the present invention.

It is possible to employ parallel processing units to carry out parts of the genetic algorithm. For example, as described in our co-pending United Kingdom patent application no. [agent's ref: HL 82083], the scheduling apparatus may have a processing unit per individual which carries out a mutation process and calculates a fitness of the individual. Other processes of the genetic algorithm, such as selection and crossover, are carried out by a central command unit which is connected to all the individual processing units.

CLAIMS:

1. A packet scheduling method, for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which
5 method comprises:

generating a plurality of candidate scheduling solutions, at least one said candidate solution being generated using a genetic algorithm, and each candidate solution specifying at least the receiver(s) to which
10 packets are to be transmitted in a scheduling instant under consideration; and

comparing the generated candidate solutions and, based on the comparison results, selecting a best one of the candidate solutions to use to transmit packets
15 in said scheduling instant.

2. A method as claimed in claim 1, wherein said genetic algorithm operates over a series of iterations and has a plurality of individuals, each representing one of said plurality of candidate scheduling
20 solutions, and in each said iteration parent individuals of a current generation produce child individuals of a next generation according to said genetic algorithm.

3. A method as claimed in claim 2, wherein each
25 said individual has a chromosome comprising genes representing respectively parameters of said candidate scheduling solution, and the genetic algorithm processes the genes of said parent individuals of said current generation to produce the genes of said child
30 individuals of said next generation.

4. A method as claimed in claim 2 or 3, wherein in a first generation at least one of the individuals represents a candidate solution generated using a deterministic scheduling method.

35 5. A method as claimed in claim 4, wherein one said individual in said first generation represents a

round-robin candidate scheduling solution generated by a round-robin scheduling method.

6. A method as claimed in claim 4 or 5, wherein one said individual in said first generation represents 5 a maximum carrier-to-interference ratio candidate solution generated by a maximum carrier-to-interference ratio scheduling method.

7. A method of as claimed in any one of claims 2 to 6, wherein in a first generation at least one of the 10 individuals represents a randomly- or pseudorandomly-generated candidate solution.

8. A method as claimed in any preceding claim, wherein a plurality of channels are available for transmitting packets from the transmitter to the 15 receivers in each scheduling instant, and each said candidate solution further specifies one or more parameters applicable to each said available channel in the scheduling instant under consideration.

9. A method as claimed in claim 8, wherein each 20 said candidate solution specifies how the specified receiver(s) is (are) to be allocated to the available channels in the scheduling instant under consideration.

10. A method as claimed in claim 9, when read as appended to claim 3, wherein each gene corresponds to 25 one of said available channels and specifies the receiver allocated to its corresponding channel in the scheduling instant under consideration.

11. A method as claimed in any one of claims 8 to 10, wherein each said candidate solution specifies a 30 modulation and/or coding scheme to be applied to each said available channel in the scheduling instant under consideration.

10. A method as claimed in any one of claims 8 to 11, wherein each said candidate solution specifies a 35 transmission power for each said available channel in the scheduling instant under consideration.

13. A method as claimed in any one of claims 8 to 12, wherein each said candidate solution further specifies an amount of data to be transmitted via each said available channel in the scheduling instant under 5 consideration.

14. A method as claimed in any preceding claim, wherein, for each said candidate solution, at least one fitness measure is determined, the or each said fitness measure being a measure of performance of the candidate 10 solution concerned.

15. A method as claimed in claim 14, wherein a set of two or more such individual fitness measures is determined for each said candidate solution, which fitness measures relate to different respective 15 performance aspects.

16. A method as claimed in claim 15, wherein said two or more individual fitness measures in said set for each candidate solution are combined together using a fitness function to produce a unified measure of 20 fitness for the candidate solution concerned.

17. A method as claimed in claim 16, wherein said genetic algorithm is operable to rank the individuals of a generation based on the respective unified fitness measures for the candidate solutions represented by 25 these individuals.

18. A method as claimed in claim 16 or 17, wherein at least one said individual fitness measure of said set is adjusted to reduce a mismatch between it and another one of the individual fitness measures of 30 said set.

19. A method as claimed in claim 16, 17 or 18, wherein, for one or more said individual fitness measures of said set, a mapping function is applied to the fitness measure so as to reduce a mismatch between 35 it and another one of said individual fitness measures of said set.

(20. A method as claimed in any one of claims 16
to 19, wherein one or more of the individual fitness
measures of said set are normalised.

5 21. A method as claimed in any one of claims 16
to 20, wherein one or more of the individual fitness
measures of said set are weighted when the individual
fitness measures are combined together to produce said
unified measure of fitness.

10 22. A method as claimed in any one of claims 16
to 21, wherein said unified measure of fitness is
produced by forming a product of said individual
fitness measures in said set.

15 23. A method as claimed in any one of claims 14
to 22, wherein one said fitness measure is a measure of
total data throughput to the specified receiver(s) for
the candidate solution concerned.

20 24. A method as claimed in any one of claims 14
to 23, wherein one said fitness measure is a measure of
an amount of data at the transmitter waiting for
transmission to the specified receiver(s) for the
candidate solution concerned.

25 25. A method as claimed in any one of claims 14
to 24, wherein one said fitness measure is a measure of
delay(s) in transmitting data to the specified
receiver(s) in the candidate solution concerned.

30 26. A method as claimed in any one of claims 14
to 25, wherein one said fitness measure is a measure of
how said candidate solution affects a fairness of said
packet scheduling amongst active ones of said plurality
of receivers.

27. A method as claimed in claim 25, wherein the
or each said delay is calculated relative to a
preselected delay parameter applicable to a type of
data to be transmitted to the receiver concerned.

35 28. A method as claimed in claim 27, wherein the
delay parameter applicable to a higher-priority type of

data is lower than the delay parameter applicable to a lower-priority type of data.

29. A method as claimed in claim 23, wherein said measure of total throughput is calculated by estimating 5 a total amount of data which will be transmitted successfully to the specified receiver(s) of the candidate solution in the scheduling instant under consideration.

30. A method as claimed in claim 24, wherein said 10 measure of waiting data is a ratio of a current amount of data waiting for transmission at said transmitter to each specified receiver to a total amount of data which has arrived at said transmitter for transmission to the receiver concerned.

31. A method as claimed in claim 25, wherein said 15 measure of delay is calculated based on a sum of respective delay times for transmissions to the specified receiver(s), each said delay time being dependent upon a difference between a current time and 20 a time of arrival at said transmitter of the earliest data currently waiting for transmission to the specified receiver.

32. A method as claimed in claim 26, wherein said 25 fairness measure is based on a variance or average of respective estimated data throughputs to all active ones of said receivers if said candidate solution is used.

33. A method as claimed in claim 17, wherein the 30 individuals of a generation are selected as parent individuals in dependence upon their rankings.

34. A method as claimed in claim 2 or 3, wherein the genetic algorithm causes two parent individuals to have two child individuals and the genes of the child individuals are dependent on the genes of their parent 35 individuals.

35. A method as claimed in claim 34, wherein the parent-individual chromosomes are subject to crossover and/or mutation processes to generate the child-individual chromosomes.

5 36. A method as claimed in claim 35, wherein in said crossover process at least one subset of the genes of each chromosome is defined and at least one gene in said subset in said first child individual is derived from a corresponding gene in said subset in said second parent individual and at least one gene in said subset in said second child individual is derived from a corresponding gene in said subset in said first parent individual.

10 37. A method as claimed in claim 36, wherein at least one gene not in said subset in said first child individual is derived from a corresponding gene not in said subset in said first parent individual, and at least one gene not in said subset in said second child individual is derived from a corresponding gene not in said subset in said second parent individual.

15 38. A method as claimed in claim 36 or 37, wherein the or each said subset, if any, is selected randomly or pseudorandomly.

20 39. A method as claimed in any one of claims 34 to 38, wherein in said mutation process at least one gene in at least one child individual is subject to a change.

25 40. A method as claimed in claim 39, wherein said gene is changed to specify a receiver having data waiting at said transmitter.

30 41. A method as claimed in any preceding claim, wherein said transmission is a wireless transmission.

35 42. A method as claimed in any preceding claim, wherein said transmitter is part of a base station of a wireless communication system, and each said receiver is part of a user equipment of that system.

43. Packet scheduling apparatus, for scheduling packets of data for transmission from a transmitter to a plurality of receivers via at least one channel, which apparatus comprises:

5 solution generating means for generating a plurality of candidate scheduling solutions, at least one said candidate solution being generated using a genetic algorithm, and each said candidate solution specifying at least the receiver(s) to which packets
10 are to be transmitted in a scheduling instant under consideration; and

best solution selecting means for comparing the generated candidate solutions and selecting, based on the comparison results, a best one of the candidate
15 solutions to use for transmitting packets in the scheduling instant.

44. A transmitter comprising:

packet scheduling apparatus as claimed in claim
43; and
20 transmitting means connected operatively to said packet scheduling apparatus and operable, when said scheduling instant occurs, to cause packets to be transmitted to the receiver(s) specified by the best solution selected for that instant by the packet scheduling apparatus.
25

45. A packet scheduling method, for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which method comprises:

30 generating a plurality of candidate scheduling solutions, each candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration;

35 producing a measure of fitness for each said candidate solution, said fitness measure being a measure of performance of the candidate solution in at

least one of the following aspects: delay, fairness, and quality of service; and

employing the measures of fitness produced for different candidate solutions to make a selection amongst the generated candidate solutions.

46. A packet scheduling method, for scheduling packets of data for transmission from a transmitter via at least one channel to a plurality of receivers, which method comprises:

generating a plurality of candidate scheduling solutions, each candidate solution specifying at least the receiver(s) to which packets are to be transmitted in a scheduling instant under consideration;

producing a measure of fitness for each said candidate solution, said fitness measure being a measure of performance of the candidate solution in at least two different aspects; and

employing the respective measures of fitness produced for different candidate solutions to make a selection amongst the generated candidate solutions.

47. A method as claimed in claim 46, wherein said different performance aspects include two or more of the following: throughput, delay, fairness, and quality of service.

48. A method as claimed in claim 46 or 47, wherein in producing said fitness measure a weighting of at least one said performance aspect is adjustable relative to that of another said performance aspect.

49. A method as claimed in claim 48, wherein each said performance aspect has an individually-adjustable weighting.

50. A method as claimed in any one of claims 45 to 49, wherein the fitness measures are employed to select a best one of the candidate solutions to use to transmit packets to the receivers in the scheduling instant under consideration.

51. A method as claimed in any one of claims 45
to 50, wherein at least one of the candidate solutions
is generated using a genetic algorithm and the fitness
measures are employed in a selection process of the
5 genetic algorithm.

52. A packet scheduling method substantially as
hereinbefore described with reference to any of Figs. 4
to 27 of the accompanying drawings.

53. Packet scheduling apparatus substantially as
10 hereinbefore described with reference to any of Figs. 4
to 27 of the accompanying drawings.

54. A transmitter substantially as hereinbefore
described with reference to any of Figs. 4 to 27 of the
accompanying drawings.

MISSOURI BLANK FORM



Application No: GB 0216245.1
Claims searched: 1 to 54

Examiner: Daniel Voisey
Date of search: 16 January 2003

Patents Act 1977 : Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance	
A		GB 2361147 A	(MOTOROLA) see particularly page 3 line 30 to page 5 line 21, page 7 line 8 to page 8 line 2 and page 19 line 22 to page 20 line 11.
A		GB 2313264 A	(HARRIS) see particularly page 2 line 21 to page 3 line 19.
A		GB 2299729 A	(NORTHERN TELECOM) see particularly page 1 line 19 to page 2 line 22.
A		GB 2355623 A	(ERICSSON) see abstract.

Categories:

- | | |
|---|--|
| X Document indicating lack of novelty or inventive step | A Document indicating technological background and/or state of the art |
| Y Document indicating lack of inventive step if combined with one or more other documents of same category. | P Document published on or after the declared priority date but before the filing date of this invention. |
| & Member of the same patent family | E Patent document published on or after, but with priority date earlier than, the filing date of this application. |

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC⁵:

H4K, H4L

Worldwide search of patent documents classified in the following areas of the IPC⁷:

H04L, H04Q

The following online and other databases have been used in the preparation of this search report:

Online: WPI, EPODOC, JAPIO and the Internet

THIS PAGE BLANK (USPTO)